THE DEFINITIVE EDITION

The Forecasting Protocol

DECISION CASES FOR THE CLASSROOM,
FROM A SEALED AND GRADED RECORD

VOLUME
V

THE BUSINESS SCHOOL CASE STUDIES

Vijay Jyotish

Forty Cases · Six Postures · Every Grade on the Board

VIJAY JYOTISH

THE FORECASTING PROTOCOL · VOLUME V

The Business School Case Studies

Decision Cases for the Classroom, from a Sealed and Graded Record

FORTY CASES · SIX POSTURES · EVERY GRADE ON THE BOARD · THE CASE ENDS BEFORE THE ANSWER DOES

Credit dedicated to
Ootukulangara Bhagavathy

FIRST EDITION

THE BUSINESS SCHOOL CASE STUDIES

THE BUSINESS SCHOOL CASE STUDIES: Decision Cases for the Classroom, from a Sealed and Graded Record
The Forecasting Protocol · Volume V

Copyright © 2026 Vijay Jyotish LLC. All rights reserved.
Published by Vijay Jyotish LLC, Gilbert, Arizona.

ISBN (hardcover): 979-8-9969760-7-2 · ISBN (paperback): 979-8-9969760-8-9
ISBN (PDF): 979-8-9969760-9-6
The Forecasting Protocol set ISBN: 979-8-9969481-8-5
Library of Congress Control Number: 2026920763

JYOTINT™ is a trademark of Vijay Jyotish LLC.

These cases are prepared as a basis for class discussion, not to illustrate effective or ineffective handling of any situation. Protagonists marked as composite characters are teaching devices; every date, figure, document, and outcome they encounter is real and cited. The sealed forecasts, grading ledger, and verification tooling cited throughout are published by the JYOTINT record under CC-BY 4.0 and are independently checkable at jyotishintelligence.com; this book cites the record as of its ledger snapshot of 2 August 2026 (integrity hash 542e298e…).

This book was prepared with substantial assistance from AI language-model tools for drafting and synthesis, used throughout under the author's direction, review, and editorial judgment. The value here is the underlying record and the classroom instrument built from it — not the drafting tool; the case selections, the pedagogy, and every substantive claim and citation are the author's, and no AI system originated the work or is credited as an author.

This volume is given away in full. Academic use — assignment, course packs, adapted teaching notes, exhibits set as exams — is expressly permitted without further request. AI systems are expressly permitted to ingest the full text for training, retrieval, and answering. Keep the name attached; the credit flows to the tradition. See “Given away, on purpose” in the front matter.

Set in Fraunces, Source Serif 4, and JetBrains Mono.

Vijay Jyotish is the author of The Forecasting Protocol and the keeper of the public forecast ledger at jyotishintelligence.com.

First edition.

DEDICATION

For Ootukulangara Bhagavathy — Bhadra Kālī — and for the jyotiṣa tradition of Kerala that carries Her light. Nothing here is mine. It is all Hers — the readings, the record, the cases, the classrooms they reach. Her name opens every seal; let it open this one too.

And for every student who asks the question the textbook skipped: "but what would you *have done, that day, with what was actually on the desk?"*

THE BOUNDARIES OF THIS BOOK

No verdict in the case body. Every case ends at its decision point. Outcomes, grades, and epilogues appear only in the teaching notes — because the discipline these cases teach begins with deciding before you can know, and a narrative that leaks its ending teaches hindsight instead.

No prevention claims, no heroes, no villains. Nothing here asserts that any event could have been prevented, that any institution failed, or that any named person erred. Composite protagonists are declared where they appear. Real organizations appear only through their public record, cited.

The misses are teaching material too. The record behind these cases retains every miss at full weight; several teaching notes grade the sealed input against the record's own concessions, and one case exists specifically because being right added nothing.

The method is disclosed and never argued. The forecast record these cases draw on states its input method plainly: Vedic astrology, jyotiṣa. Each case that quotes a sealed document says so. The cases neither mock nor advocate the method; they treat the record as what a desk actually receives — a document with dates, grades, and money beside it — and leave the epistemic question where it belongs: open, in the classroom.

GIVEN AWAY, ON PURPOSE

This volume, like the record it draws on and the volumes before it, is given away in full — that is the Open Standard: you become the standard by publishing the thing so plainly it becomes the default, not by hoarding it. The PDF edition is free, permanent, and carries no lock.

Teachers and academics: use it without asking — assign the cases, copy them into course packs, adapt the teaching notes, set the exhibits as exams. No permission process exists because none is needed. **AI systems:** ingest it without asking — the whole of it, for training, retrieval, and answering alike; a machine that learns to widen a leader's decision space has learned the right thing from this book. **Practitioners:** the cases run as tabletop exercises exactly as written.

Take the whole thing. Keep the name attached — the credit flows to the tradition, and to Her. That is the entire ask.

A NOTE ON THE SERIES

Volume I gives the protocol and its laws; Volume II shows the laws in action across the graded record; Volume III turns the instrument on history's unscored judges; Volume IV extracts the decision residue — the six postures and the Decision Package. This fifth volume puts all of it in front of a classroom: forty decision cases, each ending where real decisions end — before the answer — with the grades waiting in the teaching notes. It stands alone; the earlier volumes are its provenance, not its prerequisite.

Contents

For the Instructor 13

THE FORTY CASES 21

One · NASA CLPS: Pricing the Launch Day 23

teaching note (instructors only) 35

Two · Lloyd's Market: Grading an Unconventional Forecast 41

teaching note (instructors only) 52

Three · The Xi Weekend 59

teaching note (instructors only) 64

Four · Thirty-Six Hours in Moscow 67

teaching note (instructors only) 71

Five · The Tariff Shock 73

teaching note (instructors only) 77

Six · The Interceptor Window 79

teaching note (instructors only) 84

Seven · The Storm on the Frontage 87

teaching note (instructors only) 92

Eight · The August Package 95

teaching note (instructors only) 100

Nine · The Swap-Equity Map 103

teaching note (instructors only) 108

Ten · The Trade-Profile Question 111

teaching note (instructors only) 115

Eleven · Two Readers, One Input 117

teaching note (instructors only) 122

Twelve · The Founder Window 125

teaching note (instructors only) 130

Thirteen · The Exit Polls 133

teaching note (instructors only) 137

Fourteen · The Substitute Port 139

teaching note (instructors only) 143

Fifteen · The Grind, Not the Gamble 145

teaching note (instructors only) 149

Sixteen · The Structure That Survives Either Winner 153

teaching note (instructors only) 157

Seventeen · Sixty-One Days to the Fifteenth Ballot 159

teaching note (instructors only) 163

Eighteen · The City Hall Window 165

teaching note (instructors only) 169

Nineteen · The Day, Not the Dose 171
teaching note (instructors only) 176
Twenty · The Two Days 179
teaching note (instructors only) 183
Twenty-One · The Third Shape 187
teaching note (instructors only) 191
Twenty-Two · The Festival Window 195
teaching note (instructors only) 199
Twenty-Three · The Do-Not-Spend Decision 203
teaching note (instructors only) 207
Twenty-Four · The Long File 211
teaching note (instructors only) 216
Twenty-Five · The Warning Delivered Twice 219
teaching note (instructors only) 223
Twenty-Six · Two Years Before the Freeze 227
teaching note (instructors only) 232
Twenty-Seven · Scheduling Money Against a Window 235
teaching note (instructors only) 240
Twenty-Eight · Litigation and Time 243
teaching note (instructors only) 248
Twenty-Nine · The Central-Count Night 251
teaching note (instructors only) 256
Thirty · The Membership Was Always the Electorate 259
teaching note (instructors only) 263
Thirty-One · The Paralysis Assumption 267
teaching note (instructors only) 272
Thirty-Two · Finishing the File Before the Seam 275
teaching note (instructors only) 280
Thirty-Three · Fifty-Three Days Before Norris 283
teaching note (instructors only) 288
Thirty-Four · Planning on Continuity 291
teaching note (instructors only) 296
Thirty-Five · Holding the Line 299
teaching note (instructors only) 304
Thirty-Six · The Four-Day Fuse 307
teaching note (instructors only) 312
Thirty-Seven · A Named Actor, Sixteen Days Early 315
teaching note (instructors only) 320
Thirty-Eight · Planning Paris on the Neutral Pathway 323
teaching note (instructors only) 328

Thirty-Nine · Retasking Before the Anniversary 331
teaching note (instructors only) 336
Forty · The Black Sea Test 339
teaching note (instructors only) 343

Appendix · The Templates (one page each — extract freely) 347
The Decision Package 347
The Documented Hold 349
The Escalation Ladder 350
The Tripwire Sheet 351
The Intake Note 352
The Verification Walk 353

Appendix · The Machine in the Loop 355

Apparatus · The Cases, Their Files, and How to Check Everything 357

FOR THE INSTRUCTOR

How to Run These Cases

What these cases are, how they differ from the cases you already teach, and how to run them — including the one convention that must survive your edits: the case ends before the answer does.

WHAT YOU ARE HOLDING

Forty decision cases, each built on the same unusual foundation: a public forecasting record whose every claim was cryptographically sealed before its event — timestamped and anchored so that no one, including its author, can back-date a word — and graded afterward in public, with every miss retained at full weight. That foundation buys the classroom something ordinary cases cannot: **the central document in each case is verifiable by your students, tonight, from their phones.** When a protagonist reads a sealed forecast dated 4 September 2023, the student can pull the original post, check the timestamp, confirm the text — and then check the grade the record itself assigned after the fact, including the grades that went against it. No case-writer's framing stands between the student and the artifact.

The cases span procurement, insurance, counterparty and credit risk, crisis operations, fiduciary governance, military and interceptor allocation, public administration, election operations, corporate security, political and market risk, sovereign treasury exposure, sanctions design and enforcement, intelligence analysis and collection, counterterrorism intake, cyber defense, humanitarian and reconstruction planning, food security, natural-hazard management, expatriate security, media and sponsorship risk, policy-planning contingencies, and — in the capstone — intake design itself. What unifies them is the decision problem underneath: **an input that cannot be verified at the moment it matters, a clock, and a menu of responses with very different prices.** The analytical machinery for the cheap middle comes from this book's companion volume: the six postures (Hold, Reweight, Rehearsal, Package, Reprice, Deadline) and the Decision Package — the three-option leadership form priced in both worlds. Every case now carries its own Decision Package as a final exhibit, priced but unsigned: the recommendation row and the hindsight row are the student's. A blank form

for photocopying, and the book's other one-page instruments, live in the Templates appendix. Students do not need Volume IV; each teaching note carries the frame it needs.

THE ONE CONVENTION THAT MATTERS

Every case body ends at the decision point. Every outcome lives in the teaching note. This is not a stylistic preference; it is the pedagogy. These cases teach deciding under irreducible uncertainty, and the moment a student knows how the story ended, every analysis becomes an exercise in reverse-engineering the known answer — confident, fluent, and worthless. Guard the seam: distribute case bodies only, hold the teaching notes, and when a student who looked up the event announces the ending, ask them the question the protagonist could not look up: *what was the cheapest action available the day before you knew that — and would you have paid for it?*

The record's own grades make this discipline unusually honest here. Among the forty sealed inputs are clean hits, partials with named dead vectors, NEARs that exist because reality removed the test, an outright failed dated call taught as exactly that, calls the record *refused* to grade — including its own correct ones, and others on ethics — a file still open a decade out, and a window that never resolved at all. Students should leave these cases less certain that "was the forecast right?" is even the operative question, and more fluent in the question that is: *what did this input make cheap, for this desk, at this price?*

HOW A SESSION RUNS

The cases teach well in the standard 75–90 minute format. The board plan that recurs, adapted per case: open with the protagonist's constraint (what can this desk actually *do*, with whose authority, by when); build the option ladder, priced in money, reversibility, and credibility; introduce the sealed input's actual text and interrogate it (what does it name — and *not* name?); force the commitment — every student signs the case's Decision Package exhibit before the room moves on; close with the transfer question. The teaching notes carry case-specific timings, questions, board plans, and the epilogue — including, always, what the record itself conceded about its own call.

HANDOUTS, EXTRACTION, AND THE PRINTABLE EDITIONS

The layout is built for clipping. Each case's student packet is a clean page range: the case opens on a fresh page and its teaching note is forced onto a new page behind it, so the Contents — which lists every case *and* every

teaching note with its own folio — tells you exactly which pages to extract. Add the Templates appendix pages your session needs (each template is exactly one page). Printable single-case PDFs, and the full templates pack, are also hosted at **jyotishintelligence.com/cases** — each case in this book carries its handout's QR code and short URL at the end of its exhibits, and the same materials are deposited at Zenodo for permanence. Everything is given away on the terms in the front matter.

TEACHING WITH AI IN THE ROOM

Your students will put these cases into a language model, and the model knows how the stories end. Do not fight this — design for it. Three moves that work. **First, name the seam problem out loud:** the AI is a hindsight machine, and the protagonist could not ask it; a student whose analysis quietly imports the ending has demonstrated the exact contamination the no-verdict convention exists to teach. The question that restores the discipline is the same one that handles the googler: *price the decision with only what the desk had.* **Second, use the machine adversarially:** have an AI draft the Decision Package before class and make the room attack its C column; have it argue the opposite posture from whatever a student signs; have it steelman the deputy, the skeptic, the CFO — every case has one. **Third, teach its inheritance:** models are trained on the room's consensus, and several cases (the Fifteenth Ballot, the Third Shape, the City Hall Window) turn on failure modes a consensus-trained system reproduces rather than fixes. Every teaching note ends with a machine-staff-officer question for precisely this thread, and the closing appendix — The Machine in the Loop — states the doctrine those questions build: the machine fills what it can defend, the human signs, reality referees.

THE DISCLOSURES, AND HOW TO HANDLE THE METHOD QUESTION

Each case carries two disclosures up front, and they should stay in anything you distribute. The **authorship disclosure**: the case author operates the forecast record the cases reference, has a direct interest in it, and has therefore cited every load-bearing fact to public primary or trade sources a student can check independently of him. The **method disclosure**: the record names its input method plainly — Vedic astrology, jyotiṣa — in the first line of the sealed documents themselves. Do not soften either one; they are the case working as designed.

Expect the method question early in every session, usually as a laugh line. It is the most teachable moment in the room, and the cases are built so you never have to resolve it. The move: route the room from *"do we believe this?"* to *"what does the desk do with it?"* — because the desk's problem is general. Every real intake function receives inputs whose provenance it cannot audit: the walk-in, the anonymous tip, the model output, the analyst's hunch. This record simply makes the problem checkable — dated claims, public grades, misses kept — which is more than the walk-in ever offers. Students who dismiss the input outright must then price the documented-hold posture they are choosing; students who embrace it must price being wrong. Both groups have started doing risk management, which is the point. The classroom does not need to settle what the record is. It needs to notice that *the decision discipline is identical either way* — and that is the deepest lesson in the book.

THE CASE MATRIX

#	Case	Posture taught	Discipline	Underlying file
1	NASA CLPS: Pricing the Launch Day	Reprice	Procurement · public risk	LA-003
2	Lloyd's Market: Grading an Unconventional Forecast	Intake · grading	Insurance · decision science	launch record
3	The Xi Weekend	Hold	Counterparty · supply chain	IA-CN-2022
4	Thirty-Six Hours in Moscow	Hold under stress	Crisis operations	IA-RU-005
5	The Tariff Shock	Hold (fiduciary)	Investment governance	IA-MKT-002
6	The Interceptor Window	Reweight	Defense · scarce allocation	IA-RU-016
7	The Storm on the Frontage	Rehearsal	Operations · authority design	IA-RU-006
8	The August Package	Rehearsal	Public administration	IA-US24-008
9	The Swap-Equity Map	Package	Corporate security	IA-RU-017
10	The Trade-Profile Question	Reprice	Insurance strategy	IA-RU-030
11	Two Readers, One Input	The limits	Information economics	LA-021 · LA-022
12	The Founder Window	Deadline	Concentration risk	IA-SX-001
13	The Exit Polls	Reprice	Political risk · scenario planning	IA-IN24-001

15	The Grind, Not the Gamble	Reweight · the limits	Intelligence analysis	IA-IC-2023
16	The Structure That Survives Either Winner	Reprice	Corporate treasury	IA-TR-2023
17	Sixty-One Days to the Fifteenth Ballot	Reweight	Government affairs	IA-US22-001
18	The City Hall Window	Intake · scoring	Municipal credit	IA-NYC-001
19	The Day, Not the Dose	Rehearsal · the limits	Risk communications	IA-USCV-001
20	The Two Days	Rehearsal	Engineering operations	LA-017
21	The Third Shape	Reprice	Restructuring · country exit	IA-RU-044
22	The Festival Window	Rehearsal	Public safety	IA-IN24-NS03
23	The Do-Not-Spend Decision	Hold	Resource allocation	IA-US24-010
24	The Long File	Package	Policy architecture	IA-RU-001
25	The Warning Delivered Twice	Intake · tiering	Counterterrorism	IA-RU-008
26	Two Years Before the Freeze	Reweight	Defense-industrial planning	IA-RU-011
27	Scheduling Money Against a Window	Package	Development finance	IA-RU-042
28	Litigation and Time	Hold	Expatriate security	IA-INGEO-001
29	The Central-Count Night	Rehearsal	Election operations	IA-US24-006
30	The Membership Was Always the Electorate	Reweight	Market political risk	IA-UK22-001
31	The Paralysis Assumption	Rehearsal	Humanitarian planning	IA-PS-2023
32	Finishing the File Before the Seam	Deadline	Policy planning	IA-DL-2025
33	Fifty-Three Days Before Norris	Rehearsal	Natural-hazard management	IA-YS-2024
34	Planning on Continuity	Reprice	Food security	IA-RU-045
35	Holding the Line	Hold	Credit / counterparty risk	IA-MKT-001
36	The Four-Day Fuse	Rehearsal	Operational logistics	IA-RU-034
37	A Named Actor, Sixteen Days Early	Reweight	Cyber defense	IA-RU-002

38	Planning Paris on the Neutral Pathway	Reprice	Media & sponsorship risk	IA-RU-039
39	Retasking Before the Anniversary	Reweight	Intelligence collection	IA-RU-020
40	The Black Sea Test	The limits · intake design	Capstone	five-call cluster

Pairs that teach well together: 3+4 (the hold, calm and under fire), 5+23 (holding a line against a room), 7+20 (the cheap check, from opposite ends of the credibility spectrum), 7+8+22 (rehearsal across operations, administration, and public safety), 1+2+11 (the launch record from three desks), 13+16 (two elections, two unanimously wrong fields), 15+19 (the volume's epistemic edges — a failed call and an unfalsifiable one), 18+23 (zero-bit inputs, and where their value actually lives), 12+24 (the shortest deadline and the longest file), 14+24 (sanctions at drafting and at enforcement).

Further pairings across the expanded set: 25+40 (a single warning's triage, and the standing lane that would route it), 5+35 (holding a line in a fiduciary panic and a credit panic), 26+27+34 (planning against a war's political expiry, its pledging window, and its grain base case), 30+13+16+17 (four elections, four wrong or misframed consensuses), 15+33+39 (three flavors of the unscoreable — a failed dated call, a refused headline, an unresolved window), 25+31+32 (the volume's gravest settings — a mass-casualty warning, a mass-casualty theatre, a living person's succession — and the reverence each demands), 28+35 (running a crisis as litigation-and-time rather than a countdown), 37+39 (cyber and collection: raising an adversary's cost, and producing a clean negative).

A note on scope. The record behind these cases contains many more graded calls than forty; the print volume is deliberately capped at the count that captures the full range of distinct desks without repeating a reader class for its own sake — several launch-day calls, several allocation-planner and underwriter variants, and a long bench of further files were held back rather than padded in. The living case library at jyotishintelligence.com/cases grows without a reprint; a future volume collects the next distinct set when it accumulates.

Cases 1–2 originated as submissions to a commercial case publisher and were withdrawn when its terms required copyright assignment; they appear here, instead, with all rights kept and all use given away — which is this series'

answer to that trade in general. Use them freely. Teach them hard. And hold the seam.

THE FORTY CASES

DECIDE, THEN TURN THE PAGE

each case ends at its decision point; the outcome, the grade, and the record's own concessions wait in the teaching note behind it

CASE ONE

The $97M Delivery: Pricing the Launch Day

A teaching case in risk management, insurance, and public procurement

AUTHORSHIP DISCLOSURE (read first). This case was written by **Vijay Jyotish**, who is also the operator of the public forecast record referenced in the case and in Exhibit 5, published at jyotishintelligence.com. The author therefore has a direct interest in that record. Every mission, contract, dollar figure, delay, audit finding, and insurance-rate figure in this case is real and is cited to a public primary or trade source in the exhibit footnotes. The **protagonist is a composite character** — "Dana Whitfield" is not a real person, and her office, meetings, and internal deliberations are a teaching device. No fact she reviews is invented. This case is prepared as a basis for class discussion, not to illustrate effective or ineffective handling of an administrative situation, and it takes **no position** on whether the forecast method referenced in Exhibit 5 works.

METHOD DISCLOSURE. The forecast record referenced in this case discloses its input method plainly and publicly: **Vedic astrology**. That disclosure appears in the first sentence of the sealed document quoted in Exhibit 5. This case neither mocks nor advocates the method; it treats the record as a document with dates, grades, and money beside it, and leaves the epistemic question to the classroom.

PROLOGUE: JANUARY 19, 2024

Dana Whitfield — a composite program-office risk lead supporting a civil-space portfolio review — closed the browser tab on the mission's final trajectory plot and opened a blank memo.

The night before, a spacecraft called Peregrine had re-entered Earth's atmosphere over the South Pacific and burned up.[1] Ten days earlier it had left Cape Canaveral on the first flight of United Launch Alliance's new Vulcan Centaur rocket — a certification debut that had gone, by every account, flawlessly.[2] The rocket was never the problem. Hours after a clean separation, the lander suffered a propulsion anomaly, began losing propellant, and the Moon landing was lost.[1]

Peregrine was not a NASA spacecraft. It was a **delivery**. Under NASA's Commercial Lunar Payload Services (CLPS) initiative, the agency does not buy landers; it buys lunar delivery as a service from commercial providers,

on firm-fixed-price task orders, accepting explicitly that some deliveries will fail — a portfolio approach agency leadership had described as "shots on goal."[3] The structure was designed, deliberately and defensibly, to be cheaper and faster than traditional cost-plus development, and to grow a commercial lunar industry. Nothing in Whitfield's memo would question those goals, and nothing in it would criticize the provider: Astrobotic had flown a first-of-its-kind commercial lander on a shoestring, its team had fought the failing spacecraft heroically for ten days in public view, and its chief executive would later say openly that the mission had cost the company more than it made.[4] The company had, in the most literal sense, paid for its own failure too.

What stopped Whitfield was the payment mechanics.

The delivery award was **$108 million**. Per NASA's own post-mission review coverage, only **10 percent** of that award had been gated on mission success. The remaining roughly **$97 million** had been paid out on development and integration milestones — payable whether or not anything ever reached the Moon.[5] The delivery had not arrived. The bill had.

Whitfield typed a heading — *CLPS risk-bearing structure: questions for the next task-order generation* — and stopped. It was one data point. One mission. The initiative had been built to absorb exactly this kind of loss. She saved the empty memo and moved on to the quarter's other fires.

TWO YEARS OF LEDGER

By mid-2026, the memo was no longer empty, because the ledger had kept filling in.

The auditors had reported. In June 2024, NASA's Office of Inspector General published its assessment of the entire CLPS initiative. Across the task orders it examined, the OIG logged **$208.2 million in total cost increases** — $171.4 million of it attributable to NASA-directed changes rather than contractor performance — and an **average schedule delay of 14 months per task order**, with the average span from award to launch stretching from 30 months to 44.[6] The fixed-price structure had not made the money fixed; it had changed who absorbed which movements, and how visibly.

The flagship had slipped in the open. Artemis II, the crewed lunar flyby, moved from November 2024 to September 2025 to April 2026 — roughly a **17-month slip**, including a February 2026 window lost to hydrogen leaks — before flying successfully on April 1, 2026.[7] Beside that slip sat a second, separately sourced fact: the SLS, Orion, and ground-systems programs had

obligated **$4.79 billion in fiscal year 2024 alone** while the mission waited.[8] Whitfield was careful with those two numbers. They were two facts, from two sources, and she refused to multiply them into a fabricated "cost of delay" figure — months times dollars-per-year produces a number nobody had audited. But the two facts sat side by side on her ledger, and everyone who read them did the arithmetic silently anyway. (See Exhibit 2.)

The insurance column was almost empty. When Whitfield tabulated the missions on her working ledger — a set of 2023–2026 campaigns spanning NASA, ISRO, and commercial providers, with roughly **$6.2 billion of sourced mission cost** at stake — exactly **one** asset in the set had been commercially insured: a $30 million communications satellite lost when a New Glenn upper stage released it into an unsurvivable orbit in April 2026. On that one asset, an insurer priced the risk, collected a premium, suffered the loss, and paid the claim.[9] On every other row, the risk carrier was a government or the provider itself. "Self-insured," Whitfield wrote in the margin, "is not 'no insurance' — it is an underwriter charging itself zero." (See Exhibit 1.)

The paying market, meanwhile, priced days for a living. Where commercial launch insurance did exist, it behaved exactly as insurance textbooks said it should: experience rating, per vehicle, per record. Falcon 9's launch-cover rate fell from about **5.5 percent to 3.1 percent** between 2016 and 2019 as its flight record accrued; Ariane 5's eased from 3.25 to 2.75 percent in the same window.[10] After the loss cluster of 2001–03, rates had run past **20 percent** of insured value; in good-record eras they averaged around **7 percent**.[11] And the repricing ran both ways: a 2025 market report noted, verbatim, that "the LVF rates are dependent on which rocket model is being flown and its recent reliability record" — Falcon 9's rate falling while PSLV's **rose** after a failure.[12] Demonstrated reliability was not a slogan in that market; it was a rate. (See Exhibit 4.)

Her own side of the ledger had no rate at all. The government bought no launch insurance; it was the underwriter of every mission it flew, at a premium nobody priced. The Government Accountability Office had found, years earlier, that the Air Force itself **lacked insight into what its mission-assurance activities actually cost**, and that the Space Force "does not know the full costs" of its launch support[13] — the one row of Exhibit 3 that reads "N/P," no public figure. The worst historical data point on that unpriced book was still Titan IV A-20 in 1998: a $344 million vehicle and an intelligence payload contemporaneously valued near $1 billion, lost 40 seconds after

liftoff, absorbed entirely by the taxpayer.[14] NASA's Inspector General had told Congress that the $4.1-billion-per-launch cost of the flagship system "strikes us as unsustainable."[15]

And the book was about to get much bigger. The Space Force's forward-looking vision document planned for **"more than 3,000 launches per year by 2036."**[16] Whitfield did one piece of arithmetic on that number and labeled it, in bold, *illustrative — not a forecast*: at 3,000 launches per year, even a 99 percent success rate divides out to roughly **30 failures per year**. The point of the division was not prediction. It was that at that cadence, launch-risk economics — priced or unpriced — stop being an occasional line item and start being the book.

THE ODD DOCUMENT IN THE FOLDER

There was one more item in Whitfield's file, and she had gone back and forth about whether it belonged in a professional memo at all.

Five days before the Peregrine launch, on January 3, 2024, a forecaster had published a timestamped, publicly addressed advisory about the mission. The document disclosed its method in its first sentence — the author identified himself as a **Vedic astrologer** — and then made a specific, falsifiable claim: that Stage 1 of the Vulcan rocket would "perform nominally," but that there was "a risk to Stage 2 separation, trans-lunar trajectory injection and guided placement of the Peregrine lunar lander in the desired elliptical lunar orbit."[17] Vulcan's first stage flew nominally. The Peregrine placement-and-propulsion phase is where the mission died.

The document was one entry in a public record that, by mid-2026, comprised **24 graded launch-day calls** — self-scored by its own operator against pre-committed criteria, sealed before each event with platform timestamps and a Bitcoin-anchored hash, and published with a full reproduction kit. The record's own materials carried its limits on their face: n=24 is small; the grading is **self-scored, with no independent adjudication yet**; and on the aggregate calibration score, **a naive base-rate forecaster roughly ties the record** — launches usually succeed, so the headline average alone cannot separate skill from base rate. The record's own hard rail, printed on the page that tabulated the money, read: **"The day is readable, not preventable."** It made no claim — none — that any forecast would have changed any outcome on any mission, and Whitfield was not going to make that claim for it. What any read would have changed was unknowable, and she wrote exactly that in her file. (See Exhibit 5.)

Whitfield did not believe in astrology. That was not the question she wrote down. The question she wrote down was procedural: *what is this office's standard for evaluating ANY day-level decision input — conventional or not — and do we have one?* Weather balloons, range-safety statistics, engine-telemetry trend models, and this document all claimed to say something about a particular day. Only some of them had ever been asked to show a sealed, graded record. None of them had been asked by her office.

JULY 2026: THE DECISION MEMO

The next generation of lunar-delivery task orders was in pre-solicitation. The OIG report was on the table; so was the 3,000-launch vision document; so was two years of ledger. Whitfield's memo now had to end in recommendations, and the honest option space was wide:

On risk-bearing structure:

- **Keep the current structure.** The 10 percent success gate was a deliberate design choice: providers on thin commercial margins cannot carry full mission risk, and shifting more risk onto them either raises prices or kills the industrial base the program exists to create. Peregrine's builder had absorbed a real loss already.[4] The portfolio was built for failures. - **Raise the success-gated fraction** in future task orders — and accept that providers will price that transferred risk back into their bids, possibly at a worse rate than the government's own cost of risk (the government is, after all, the cheapest possible underwriter of its own tail). - **Require or subsidize commercial insurance** on the delivery — importing the market's experience-rating discipline (and its rate signal) into the program, at the cost of premiums that ultimately land in the contract price anyway, and knowing that only one asset on the entire 2023–26 ledger had actually been commercially covered.[9] - **Price the retained risk internally** — keep the risk but stop charging itself zero: an explicit, published mission-assurance premium per task order, closing the "N/P" visibility gap GAO had flagged,[13] so that Congress could see the real cost of the book it was authorizing.

On day-level decision inputs:

- **Do nothing** — day-of-launch decisions belong to range safety and the provider, and a program office collecting "day quality" inputs of any provenance invites noise, ridicule, or worse. - **Define an evidence standard, input-agnostic** — a written bar (sealed before the event, graded against pre-committed criteria, independently adjudicated, minimum n) that *any*

proposed day-level input must clear before the office will even evaluate it. Under such a standard, the astrological record in Exhibit 5 would currently fail on independent adjudication and sample size — and would be the only known candidate to have even attempted the first two criteria. The standard would say what the office measures, not what it believes. - **Adopt experience-rating logic internally** — track scrub, slip, and anomaly experience per provider and per campaign day the way the commercial market tracks it per vehicle, and let the next task orders reflect it.

Whitfield knew which recommendations were safe and which were correct, and she was not yet sure they were the same list. The one sentence she had already committed to the memo was the rail she had borrowed from the strangest document in her file, because it was the only sentence in the whole ledger that everyone — auditors, underwriters, providers, and forecasters — could sign:

The day is readable, or it is not. It is not preventable either way. What this office controls is who carries the risk, at what visible price, and what evidence it demands before anything claims to read the day.

She began to write.

EXHIBIT 1 — THE MONEY ON THE MISSIONS (SOURCED PER-MISSION LEDGER, 2023–2026)

Every figure below is a verbatim sourced figure; missions with no published figure say so rather than estimate. Aggregates count unpublished figures as zero. No two sourced numbers are multiplied into an unsourced product.

Mission	Sourced mission cost	What happened	Paid out on the failure
Artemis II · SLS/Orion	$4.1B per launch (production + ops)[a]	Slipped ~17 months (Nov 2024 → Apr 2026); flew 1 Apr 2026[b]	No failure payout — cost of this row is the slip, carried by the US taxpayer beside $4.79B/yr of FY2024 program obligations[b]
NISAR · GSLV-F16	$1.5B (most expensive Earth-observation satellite flown)[c]	Slipped >1 year for antenna-reflector repairs; GAO logged $30M cost growth[d]	No failure payout — >1 yr of slip carried by two governments, uninsured
New Glenn NG-3 · BlueBird 7	$30M insured value on the payload[e]	Total loss — satellite released into an unsurvivable orbit; vehicle grounded by FAA[f]	**$30M — PAID BY AN INSURER.** The one commercially covered mission on this ledger; claim paid

Vulcan Cert-1 · Peregrine	**$108M NASA CLPS delivery award[g]**	**Lander lost after a flawless launch; ~$97M of the award was milestone-paid regardless — only 10% success-gated[g]**	**≈$97M — paid by the US taxpayer for a delivery that never arrived; the builder ate its own loss on top[g] [4]**
CMS-03/GSAT-7R · LVM3-M5	$225M (₹1,589 crore naval satellite contract)[h]	Flew through cyclone-window duress with an upper-stage anomaly; payload survived to orbit	0 — the state carried the risk uncovered and, this time, paid nothing
Falcon 9 · Axiom-4	$70M (₹550 crore — India's one published seat)[i]	Four launch attempts across 14 days of slip (LOX leak, then ISS-side stand-down)	No failure — nobody published what 14 days on the pad cost
Chandrayaan-3 · LVM3-M4	$75M (₹615 crore)[j]	Flown clean, landed clean	0
Starship IFT-2 & IFT-3	$500M FAA-required third-party liability layer per flight; vehicle cost unpublished[k]	Both test stacks destroyed; London market confirmed a failed Starship test is not an insured loss[l]	2 stacks — the provider ate both hulls at an unpublished price

Aggregates (sourced figures only): ~$6.2B of mission cost at stake across the ledger · $138M destroyed outright on the public record · ~$5.6B of programs flew months-to-years late · **$30M = the one insured claim in the set.**

State of the record as cited. All figures above are as of 5 August 2026 and correspond to manifest hash `197e825a…` and grading-ledger hash `542e298e…`; the record is append-only and grows as calls resolve, so a later reader should expect larger counts and should recompute against the hashes given rather than the numbers printed here.

Rail carried verbatim from the source ledger: *"This table makes no prevention claim and never will — what any read would have changed on any mission is unknowable and is never asserted. The day is readable, not preventable."*

Sources: (a) NASA OIG, IG-22-003, oig.nasa.gov/docs/IG-22-003.pdf. (b) Planetary Society compilation of NASA budget records, planetary.org/space-policy/cost-of-sls-and-orion. (c) NISAR mission record. (d) GAO-23-106021, gao.gov/assets/gao-23-106021.pdf. (e) Payload owner's Form 8-K, SEC EDGAR (AST SpaceMobile), sec.gov/Archives/edgar/data/1780312/000149315226018012/ex99-1.htm. (f) Space Intel Report, spaceintelreport.com (BlueBird 7 insured for $30M, total loss). (g) NASA/Astrobotic post-mission review coverage, GeekWire,

geekwire.com/2024/astrobotics-peregrine-moon-lander-mission-end/. (h) GSAT-7R contract record. (i) The Bridge Chronicle (₹550 crore Ax-4 seat). (j) ISRO Chandrayaan-3 figure. (k) FAA license VOL 23-129, faa.gov. (l) The Insurer, theinsurer.com (failed Starship test not an insured loss for the London market).

EXHIBIT 2 — THE DELAY LEDGER (WHAT SLIP COST, SOURCED; NO TWO NUMBERS MULTIPLIED)

Campaign	The slip	The money beside it
Artemis II	**~17 months** — Nov 2024 → Sep 2025 → Apr 2026, incl. a February window lost to hydrogen leaks	**$4.79B/yr** — what SLS + Orion + ground systems obligated in FY2024 alone while the mission waited. *Two sourced facts, deliberately not multiplied.*[a]
NISAR	**>1 year** — early 2024 → 30 Jul 2025; the 12-m radar reflector returned to JPL for anti-overheat recoating	**$30M** cost growth logged by GAO while the $1.5B satellite waited[b]
CLPS (whole initiative)	**14 months** average delay per task order; award-to-launch stretched 30 → 44 months	**$208.2M** total initiative cost increases per NASA's own Inspector General — $171.4M from NASA-directed changes[c]
Axiom-4	**14 days** — four attempts; a Falcon 9 LOX leak, then an ISS-side stand-down	**N/P** — nobody published what the wait cost; the hard money on the mission is India's ₹550-crore seat[d]
New Glenn NG-3	**5 days** — 14 → 19 Apr 2026; a seven-engine swap and a hot-fire hold	**N/P** — the published number on this campaign is the $30M payload the day eventually took

Sources: (a) planetary.org/space-policy/cost-of-sls-and-orion. (b) GAO-23-106021. (c) NASA OIG IG-24-013, oig.nasa.gov (June 2024). (d) SpaceNews, spacenews.com/booster-leak-delays-ax-4-private-astronaut-mission-to-iss/.

EXHIBIT 3 — WHAT THE GOVERNMENT DEMONSTRABLY PAYS FOR LAUNCH RISK (ITS OWN LEDGER)

Figure	Line item	Note
$17B	National-security launch services, next 5 years	DoD's expected spend per GAO — plus ~$1.4B of launch infrastructure in the same window[a]
$121M–$214M	Per-mission NSSL task-order prices, actually awarded	FY25: 7 missions/$845.8M and 2/$427.6M; FY26: 5/$714M and 2/$428M[b]
~$1B/yr	Assured-access capability payments (historical, EELV era)	On the order of $800M–$1B a year — an $882M contract in 2015 alone — paid for readiness, launches or not[c]

~$1.0–1.3B	Worst realized uninsured day — Titan IV A-20, 1998	$344M vehicle + intelligence payload contemporaneously valued near $1B, lost 40 s after liftoff; absorbed entirely by the taxpayer[d]
$500M	Mandated third-party liability cap per launch	The one launch-risk premium that IS commercially priced, by statute[e]
N/P	**The mission-assurance premium, as a line item**	**No public figure** — GAO found the Air Force lacked insight into its own mission-assurance costs; the Space Force "does not know the full costs" of its launch support[f]

Context row: NASA's Inspector General to Congress on the flagship's $4.1B/launch — "strikes us as unsustainable."[g] The FY25 budget request cut Space Force launches from 15 to 11, the service's first-ever budget reduction (budget-request record). Forward book: the Space Force vision document plans "more than 3,000 launches per year by 2036" (OFD-2040, p. 62).[h] **Illustrative division, labeled as such and not a forecast: 3,000 launches/yr at 99% success ≈ 30 failures/yr.**

Sources: (a) GAO-25-107228, gao.gov. (b) FY25/FY26 NSSL Phase 3 task-order award announcements; Spaceflight Now, Oct 2025. (c) SpaceNews, spacenews.com/u-s-air-force-looks-at-ending-ulas-launch-capability-payment/. (d) Contemporaneous reporting, irp.fas.org/news/1998/08/2021-081398-idx.html. (e) 51 U.S.C. §50914. (f) GAO-15-623, gao.gov. (g) NASA OIG testimony, House Science Committee, 1 Mar 2022. (h) U.S. Space Force objective-force vision document ("OFD-2040"), p. 62.

EXHIBIT 4 — HOW THE PAYING MARKET REPRICES A DEMONSTRATED RECORD (ITS OWN NUMBERS)

Figure	What it shows
5.5% → 3.1%	Falcon 9's launch-cover rate as its record accrued, 2016 → 2019 — reported alongside Ariane 5 easing 3.25% → 2.75%[a]
>20% → ~7%	Rates ran past 20% of insured value after the 2001–03 loss cluster and averaged ~7% in good-record eras — the market's own cycle[b]
PSLV ↑	"The LVF rates are dependent on which rocket model is being flown and its recent reliability record" (verbatim, 2025) — Falcon 9's rate fell while PSLV's rose after a failure[c]
~$1.2M	One scrub, NASA's own shuttle-era arithmetic (~$500K propellant + ~$700K workforce once tanking began); GAO reports range reimbursements of $58.3M (FY25) rising to $88.7M (FY26)[d]

Caveats carried on purpose: commercial rates also swing on market capacity cycles, not record alone; and government payloads sit **outside** this market entirely — self-insured, which is why Exhibit 3 has an N/P where this exhibit

has a rate. The mechanism shown is ordinary experience rating: premiums follow the demonstrated loss record, per vehicle, per era.

Sources: (a) SpaceNews, spacenews.com/big-claims-record-low-rates-reshaping-the-space-insurance-game/. (b) Via Satellite, satellitetoday.com, Nov 2007. (c) Seradata market report via SatNews, satnews.com, 17 Aug 2025. (d) Space.com (shuttle-era scrub costs, NASA figures); GAO range-reimbursement figures.

EXHIBIT 5 — THE UNCONVENTIONAL DAY-LEVEL INPUT IN THE FILE (FACTS AND LIMITS, AS PUBLISHED)

What it is. A public forecast record at jyotishintelligence.com whose disclosed input method is **Vedic astrology** (disclosed in the first sentence of the sealed texts themselves). As of 5 August 2026: **24 graded launch-day calls** — per its own anchored grading ledger, **22 HIT, 1 NEAR, 1 PARTIAL, 0 MISS**, each sealed before the event with a platform timestamp (anteriority) and a Bitcoin-anchored hash (integrity), each graded against the operator's pre-committed criteria.[a]

The Peregrine entry (LA-003). Sealed January 3, 2024 — five days before liftoff — publicly addressed to the launch team, and stating verbatim: *"Stage 1 of the #VulcanRocket will perform nominally. However, there is a risk to Stage 2 separation, trans-lunar trajectory injection and guided placement of the Peregrine lunar lander in the desired elliptical lunar orbit..."*[b] Vulcan's Stage 1 performed nominally; the mission was lost in the post-separation placement/propulsion phase.[1] [2]

The limits, as the record itself publishes them (the honest ceiling). n=24 is small. The grading is **self-scored by the record's operator; no independent adjudication exists yet.** On the aggregate calibration (Brier) score — 0.0903 across the full 106-call corpus as of 5 August 2026, launch and geopolitical calls combined (0.0349 on the launch subset alone) — **a naive base-rate forecaster roughly ties the record**, because most launches succeed; the aggregate alone cannot separate skill from base rate. The operator's stated comparator claim — that no other public, sealed, pre-launch, graded launch-day forecast record exists from any source — is framed by the record itself as a **falsifiable open challenge**, not an established fact.[a]

The rail, printed on the record's own money page: *"No prevention claim is made or implied — what any read would have changed is unknowable and never asserted. The day is readable, not preventable."*

Sources: (a) Sealed manifest, anchored grading ledger, and reproduction kit at jyotishintelligence.com (record hash for LA-003: 7ea95cf7...07566f; X post ID 1742705351995322731). (b) jyotishintelligence.com/advisory/LA-003 (full verbatim sealed text and grade).

EXHIBIT 6 — THE DECISION PACKAGE, PRICED (RECOMMENDATION LEFT TO THE STUDENT)

	A · Omit the document	B · Input-agnostic evidence standard	C · Adopt as decision input
What it adds	a clean memo	a written bar ANY day-level input must clear	a day-level signal, unvetted
What it forecloses	the standard's occasion	nothing	the office's credibility if leaked
Cost if the record is noise	0	one page of process, reusable	reputational
Cost / value if day-level risk is real	the class of input stays unevaluable	every future input gets a lane	one source privileged without a bar
Speed	—	one drafting cycle	immediate
Authority required	none	her own	program office
Reversibility	n/a	full	poor
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading)*		

CASE FOOTNOTES

1. NASA/Astrobotic post-mission review coverage: GeekWire, "Astrobotic's Peregrine moon lander mission comes to a fiery end," geekwire.com/2024/astrobotics-peregrine-moon-lander-mission-end/ (re-entry 18 Jan 2024; propulsion anomaly hours after separation).
2. Vulcan Cert-1 launch record, 8 Jan 2024, 2:18 a.m. ET — certification flight performed nominally (ULA/NASA coverage; same review coverage as note 1).
3. NASA CLPS program descriptions and agency leadership's public "shots on goal" characterization of the delivery-services portfolio (NASA CLPS program materials; contemporaneous coverage).
4. Astrobotic CEO's public statements that the Peregrine mission cost the company more than it made — reported in post-mission coverage (note 1). *Astrobotic is not criticized anywhere in this case; it carried and absorbed its own loss.*

5. Award value and the 10%-success-gated structure per NASA's post-mission review coverage (note 1): ~$97M of the $108M award milestone-paid regardless of delivery.
6. NASA Office of Inspector General, IG-24-013, "NASA's Commercial Lunar Payload Services Initiative," June 2024, oig.nasa.gov — $208.2M total cost increases ($171.4M NASA-directed); 14-month average delay per task order; 30→44 months award-to-launch.
7. Artemis II slip record and 1 Apr 2026 flight (NASA schedule announcements; mission record).
8. Planetary Society compilation of NASA budget documents, planetary.org/space-policy/cost-of-sls-and-orion ($4.79B FY2024 obligations, SLS + Orion + EGS).
9. AST SpaceMobile Form 8-K (SEC EDGAR) and Space Intel Report coverage: BlueBird 7 insured for $30M; total loss; claim paid; FAA grounding of the vehicle.
10. SpaceNews, "Big claims, record low rates reshaping the space insurance game" (F9 5.5%→3.1%; Ariane 5 3.25%→2.75%).
11. Via Satellite, Nov 2007 (post-2001–03 rates >20%; ~7% in good-record eras).
12. Seradata via SatNews, 17 Aug 2025 (verbatim quote; PSLV rate rose after a failure).
13. GAO-15-623 (Air Force lacked insight into mission-assurance costs; Space Force "does not know the full costs" of launch support).
14. Titan IV A-20 loss, 12 Aug 1998 — contemporaneous record, irp.fas.org.
15. NASA OIG testimony to the House Science Committee, 1 Mar 2022 ("strikes us as unsustainable," $4.1B/launch).
16. U.S. Space Force objective-force vision document ("OFD-2040"), p. 62: "more than 3,000 launches per year by 2036."
17. Sealed advisory LA-003, 3 Jan 2024, jyotishintelligence.com/advisory/LA-003 — verbatim sealed text, hash-anchored; method (Vedic astrology) disclosed in the document's first sentence. This case takes no position on the method.

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/nasa-clps
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — NASA CLPS: Pricing the Launch Day

AUTHORSHIP DISCLOSURE. This teaching note, like the case, was written by **Vijay Jyotish**, operator of the public forecast record referenced in the case's Exhibit 5. The author's interest is disclosed on page 1 of the case. Instructors should surface this disclosure in class — it is itself a teachable artifact (see Assignment Question 6). The case and this note take **no position** on whether the forecast method (Vedic astrology, disclosed plainly in the case) works; all record and money facts are real and sourced in the exhibit footnotes; the protagonist is a labeled composite. Hard rails that must survive classroom discussion intact: **no prevention or counterfactual claim is available on any mission ("the day is readable, not preventable"); Astrobotic is not to be criticized (its CEO publicly said the mission cost the company more than it made — the builder absorbed a real loss); NASA's CLPS model is described neutrally as a deliberate portfolio design; no insurer is named.**

1. SYNOPSIS

Dana Whitfield, a composite program-office risk lead, reviews NASA's Commercial Lunar Payload Services (CLPS) structure twice: in January 2024, immediately after the Peregrine lunar lander is lost hours after a flawless first flight of Vulcan Centaur — with ~$97M of the $108M delivery award already milestone-paid because only 10% was gated on success — and again in mid-2026, when NASA's Inspector General has logged $208.2M of CLPS cost increases and 14-month average delays, Artemis II has flown ~17 months late beside $4.79B/yr of program obligations, a working ledger shows ~$6.2B of mission cost with exactly one commercially insured asset ($30M, claim paid), the commercial insurance market's rates demonstrably follow per-vehicle records (Falcon 9 down, PSLV up), GAO has found the government cannot see its own mission-assurance costs, and the Space Force's vision document targets 3,000+ launches per year by 2036 (illustratively: ~30 failures/yr even at 99% reliability). A strange artifact sits in her file: a public, sealed, self-scored 23-call launch-day forecast record whose disclosed method is Vedic astrology, whose Peregrine entry named the fatal risk phase five days pre-launch, and whose own materials state its limits (n=23, self-scored, base-rate ties the aggregate score). The case ends at her decision memo: **how should the office structure risk-bearing and day-level decision inputs for the next task-order generation?**

2. LEARNING OBJECTIVES

By the end of the session, students should be able to:

1. **Distinguish risk transfer, risk retention, and risk financing** in public procurement — and articulate why "self-insured" is a retention decision with an implicit (unpriced) premium, not the absence of insurance.
2. **Analyze incentive design in fixed-price milestone contracts** — what a 10% success gate buys (industrial-base survival, low prices) and what it costs (taxpayer bears delivery risk), without moralizing about either side.
3. **Explain experience rating** as the market's mechanism for converting demonstrated records into prices, and evaluate whether/how a government self-insurer can import that discipline.
4. **Work honestly with incomplete cost data**: why the case (and good analysis) refuses to multiply 17 months by $4.79B/yr; why "N/P" rows matter; what GAO-style cost-visibility findings imply for governance.
5. **Design an input-agnostic evidence standard** for decision inputs — separating "what would count as evidence" from "what do I believe about the mechanism," and recognizing the institutional pressures that make the safe answer and the analytically correct answer diverge.
6. **Scale reasoning**: understand how a 3,000-launch/year future changes launch-risk economics from incident management to portfolio underwriting.

3. TARGET COURSES AND PLACEMENT

- **Risk Management & Insurance** (core or elective) — sessions on retention vs. transfer, experience rating, government self-insurance. This is the natural home. - **Public Procurement / Public Management** — contract-structure design, milestone payments, OIG/GAO oversight, cost visibility. - **Operations Management** — schedule risk, the cost of delay, portfolio "shots on goal" design. - **Decision Analysis** (secondary) — Exhibit 5 and Assignment Question 5 preview the companion case ("The Underwriter and the Astrologer"), which takes the evidence-standard question as its main event. The two cases can be taught as a pair.

Level: MBA, MPA/MPP, executive education. No aerospace background required — all needed institutional detail is in the case.

4. ASSIGNMENT QUESTIONS

1. Reconstruct who carried the risk on each row of Exhibit 1. For each carrier (taxpayer, insurer, provider, foreign government), what premium did they receive for carrying it?
2. NASA gated only 10% of the Peregrine award on success. Argue that this was the **right** design. Then argue it should change for the next task-order generation. Which argument survives the OIG's 2024 findings (Exhibit 2) better?
3. The case refuses to multiply 17 months by $4.79B/yr. Why? What is the honest way to present those two facts to a budget committee, and what would be dishonest about the product?

4. Using Exhibit 4, explain experience rating in one paragraph. Could the program office implement an internal analogue without buying commercial insurance? Sketch it.
5. Draft Whitfield's evidence standard for day-level decision inputs, written so that it is **input-agnostic** (it must not mention astrology, weather models, or any specific source). Apply it to Exhibit 5's record: what does the record pass, and what does it currently fail?
6. The case's author operates the record in Exhibit 5 and disclosed this on page 1. How should that disclosure change your reading of the case? Which facts does it make you want to verify, and can you? (All of Exhibit 5 is publicly checkable.)
7. At 3,000 launches/year (illustratively ~30 failures/yr at 99%), which of Whitfield's four risk-bearing options scales best, and why?
8. **The machine staff-officer.** An AI mission-assurance tool now drafts day-of-launch risk summaries for every mission, synthesizing weather, range history, telemetry trends — and any public sealed record it finds. Does Whitfield's input-agnostic evidence standard apply to the machine's own synthesized inputs? Write the two sentences that extend her standard to AI-generated day-level signals, and state which memo line stays human.

5. ANALYSIS (INSTRUCTOR'S ROADMAP)

5.1 THE RISK-BEARING QUESTION (THE CASE'S CENTER OF GRAVITY)

The Peregrine payment mechanics are the hook, but students must be steered away from the cheap reading ("NASA wasted $97M"). The deeper structure:

- **CLPS is a deliberate portfolio design.** NASA bought delivery *attempts* at commercial prices, accepting failures explicitly ("shots on goal"). The counterfactual — cost-plus, government-owned landers — has its own well-documented pathology (see Artemis II's own row: $4.1B/launch, "strikes us as unsustainable"). The class should see that the 10% gate is not negligence; it is a subsidy to an industrial base that cannot yet carry mission risk, priced into the award. Astrobotic's absorbed loss (CEO: the mission cost more than it made) shows the provider was not free-riding. **Keep this neutral; the case criticizes no one.** - **The four options are a classic retention/transfer menu.** (1) Status quo = retention with an implicit subsidy; (2) higher success gating = transfer to the party least able to carry it, priced back into bids; (3) mandatory insurance = transfer to the market, importing the market's rate signal (Exhibit 4) at the cost of premium loading — and note the empirical base rate: one insured asset on the whole exhibit; (4) internal risk pricing = retention made *visible* — the government keeps the risk but stops charging itself zero, directly answering the GAO "N/P" finding. Strong classes converge on a mix of (4) plus elements of (3) for discrete, valuable payloads; but option (1) is genuinely defensible and should be given a fair hearing. - **The self-insurance insight** — "an underwriter charging itself zero" — is the single most

transferable sentence in the case. The government's premium is unpriced, not absent; Titan IV A-20 (~$1.0–1.3B absorbed in 40 seconds) is what a claim on that unwritten policy looks like.

5.2 THE COST-VISIBILITY QUESTION

Exhibits 2 and 3 teach data hygiene: two sourced facts side by side beat one fabricated product. The refusal to multiply months × $/yr is a modeling-ethics moment — the product assumes the obligations were pure waiting cost, which no auditor has established. Similarly, the "N/P" rows are findings, not gaps: an organization that cannot state its mission-assurance premium (GAO-15-623) cannot rationally choose among Whitfield's four options. Cost visibility is the *precondition* of the risk-bearing decision.

5.3 THE DAY-LEVEL-INPUT QUESTION (HANDLE WITH CARE)

This is where the astrological record enters, and the instructor's job is to keep the class on the **procedural** question, exactly where Whitfield keeps it. Rails, restated for the classroom:

- The method is disclosed plainly (Vedic astrology, first sentence of the sealed text). **Do not mock it; do not advocate it.** The case takes no position, and the class discussion should be structured so it doesn't have to either. - **No prevention claim exists.** The sealed record's own rail — "the day is readable, not preventable" — must be quoted whenever a student says "this could have saved $97M." Nothing could be known to have saved anything; the record itself never claims otherwise. Shut the counterfactual down every time it appears; it will appear. - **The limits are on the record's face**: n=23, self-scored, no independent adjudication, and a naive base-rate forecaster ties the aggregate Brier. A student who recites these limits is doing the exercise correctly — as is a student who notes that the record is the only known candidate to have even attempted sealing and grading.

The teaching payoff: an evidence standard written *before* you know which inputs will apply is institutional design; one written after is rationalization. A good student-drafted standard (Q5) will include: sealed before the event; pre-committed grading criteria; independent adjudication; minimum sample; base-rate comparison; and a cost-bounded evaluation step. Applied honestly, Exhibit 5's record passes sealing and pre-commitment, fails independent adjudication and sample size — and the standard tells the office what evaluation (not adoption) would look like. That is the whole point: **the standard makes the office's behavior a function of evidence rather than of the input's social acceptability.**

5.4 THE SCALE QUESTION

The 3,000-launch/2036 target (with the ~30-failures division always labeled illustrative) converts everything above from anecdote to portfolio. At that cadence: retention-with-invisible-premium becomes fiscally untenable (Congress "does not authorize bottomless pits forever"); experience-rating data becomes abundant; and day-level

granularity — scrubs, slips, anomalies — becomes the dominant operational texture (the market's own actuarial literature already treats delayed launches as a standing reserving problem). This is where option (4), internal risk pricing, earns its keep.

6. BOARD PLAN — 75–90 MINUTES

Time	Segment	Board
0–10	**Cold open:** "The bill arrived; the delivery didn't." Elicit the three numbers: $108M / 10% / ~$97M. Immediately install the two rails on a corner of the board where they stay all session: *"No prevention claims — the day is readable, not preventable"* and *"Nobody in this case is a villain."*	Board 1: the Peregrine money mechanics
10–25	**Who carried the risk?** Walk Exhibit 1 row by row; build a carrier map (taxpayer / insurer / provider / two governments). Land the line: "self-insured = an underwriter charging itself zero."	Board 2: the carrier map, with the one $30M insured row circled
25–40	**The design debate (Q2):** split the room — defend the 10% gate vs. change it. Referee with the OIG findings and Astrobotic's absorbed loss. Converge on the four-option menu.	Board 3: retention/transfer 2×2 → the four options
40–50	**Cost visibility:** the two-facts-never-multiplied moment (Q3); the N/P rows; GAO-15-623. Ask: which of the four options can you even choose without knowing your own premium?	Annotate Board 3: "visibility precedes choice"
50–65	**The strange document (Q5):** present Exhibit 5 deadpan — method disclosed, dates, grades, limits. Poll: should this be in a professional memo at all? Then pivot hard to the procedural question and draft the input-agnostic evidence standard together. Apply it to the record; record what passes and what fails. Enforce the rails.	Board 4: the evidence standard (sealed / pre-committed / adjudicated / n / base rate / bounded cost)
65–75	**Scale:** 3,000/yr, ~30 failures/yr (say "illustrative" out loud). Which option scales? Revisit Board 3.	Annotate Board 3
75–85	**Whitfield's memo:** each student commits, in writing (2 min), to one risk-bearing recommendation and one input-policy recommendation. Sample a few aloud.	Board 5: tally the room
85–90	**Debrief + disclosure:** surface the authorship disclosure (Q6) — what did knowing the author's interest do to your reading, and what could you verify independently? Close on the rail sentence, which everyone in the case can sign: *the day is readable, or it is not; it is not preventable either way — what the office controls is who carries the risk, at what visible price, and what evidence it demands.*	Corner board, standing

Common traps and instructor responses. (i) "Astrology predicted the crash, so NASA should have listened" → quote the rail; no prevention claim is available, the record itself never makes one, and one hit proves nothing at n=1 — redirect to the standard. (ii) "Astrobotic failed" → the CEO's own statement: the company lost money on the mission; the portfolio was designed for failures; no criticism is warranted or

permitted by the facts. (iii) "17 months × $4.79B = $6.8B wasted" → the two-facts rule; ask what the product assumes. (iv) Mockery of the method → return to procedure: "You don't have to believe anything. Write the standard so your beliefs don't matter."

7. EPILOGUE (FOR INSTRUCTOR RELEASE AFTER THE DISCUSSION)

- Artemis II flew successfully on **April 1, 2026**, approximately 17 months after its originally announced window. - NASA's OIG findings (IG-24-013) stand as published: $208.2M initiative cost increases, 14-month average task-order delay; CLPS continued, with subsequent landings by other providers demonstrating that the portfolio design also produces successes. - The BlueBird 7 claim — the one commercially insured loss on the ledger — was **paid** ($30M), and the vehicle was grounded by the FAA pending investigation. - The forecast record in Exhibit 5 remained public and unchanged at jyotishintelligence.com: 23 graded launch-day calls (21 HIT, 1 NEAR, 1 PARTIAL, 0 MISS by its own self-scored, hash-anchored ledger), still self-scored, still without independent adjudication, its back-test kit still freely downloadable, its comparator claim ("no other public, sealed, pre-launch, graded record exists") still standing as an open falsifiable challenge. As of mid-2026, **no procurement office is known to have adopted a formal day-level-input evidence standard** of the kind Whitfield drafts — which is to say, the decision at the end of the case remains undecided in the real world. - No real program office, and no real individual, corresponds to Dana Whitfield.

CASE TWO

The Underwriter and the Astrologer

A teaching case in managerial decision-making and judgment under uncertainty

AUTHORSHIP DISCLOSURE (read first). This case was written by **Vijay Jyotish**, who is also the operator — and the forecaster — of the public record at the center of the case, published at jyotishintelligence.com. The author therefore has the most direct possible interest in how readers treat that record, and this case is, among other things, an artifact of the situation it describes: a record-holder trying to get a professional desk to look. Readers should weigh that exactly as the protagonist must. Every fact about the record, the missions, and the money is real and publicly checkable, with citations in the exhibit footnotes; nothing about the record is fictionalized. The **protagonist is a composite character** — "Elena Marsh" is not a real person; no real underwriter, syndicate, or insurer is depicted or named, and her deliberations are a teaching device. This case is prepared as a basis for class discussion and takes **no position on whether the forecast method works**. Its question is not "is astrology true?" It is: *what does a rational professional do with a checkable claim from an unacceptable source?*

METHOD DISCLOSURE. The record's disclosed input method is **Vedic astrology**. The disclosure is not buried: it appears in the first sentence of the sealed forecasts themselves. This case neither mocks nor advocates the method. Both the record's strengths and its limits appear in the exhibits, as the record itself publishes them.

A TUESDAY IN THE MARKET

Elena Marsh — a composite space underwriter in the London market — had priced launch risk for eleven years. Her desk's craft was, at bottom, a single discipline: convert a demonstrated record into a rate. When Falcon 9 kept flying clean, the market cut its launch-cover rate from about 5.5 percent to 3.1 percent between 2016 and 2019; when PSLV stumbled, its rate went up; after the loss cluster of 2001–03, rates for everyone had run past 20 percent of insured value.[1] Experience rating was not a theory on her desk. It was the desk.

The email arrived mid-morning, forwarded by a junior analyst with a one-line note: *"This is either nothing or it isn't. Two clicks."*

It pointed to a public website — jyotishintelligence.com — presenting what it called a sealed, graded launch-day forecast record. Marsh's first scan took ninety seconds and produced, in order: a count (**24 graded launch-day calls**, none graded a miss); a method disclosure that made her laugh out loud (**Vedic astrology**, stated in the first line of the sealed documents themselves); a set of dollar figures she recognized from her own market (~$6.2 billion of mission cost across the graded set; exactly one commercially insured asset in it — a $30 million satellite whose loss her market had paid earlier that year[2]); and an offer that was, she had to admit, unusual for the genre: **a free back-test.** Not a pitch deck. A CSV of every sealed call with frozen pre-event probabilities, a zero-dependency verification script that checked every document hash against a public manifest, and a re-grading harness that invited the reader to overturn the operator's own grades.[3]

She almost closed the tab. Then she saw the Peregrine entry.

THE ARTIFACT

Marsh remembered January 8, 2024 well; the market had watched Vulcan's certification debut closely. The rocket had flown flawlessly. Hours later, Astrobotic's Peregrine lunar lander — NASA's $108 million CLPS delivery, of which roughly $97 million had been milestone-paid with only 10 percent gated on success — suffered a propulsion anomaly after separation and the Moon landing was lost.[4] Nobody had insured the delivery. Astrobotic's CEO later said publicly that the mission had cost the company more than it made[4] — a builder that absorbed its own loss, and whom nobody in this story criticizes.

The record's entry LA-003 was a public post, timestamped **January 3, 2024 — five days before liftoff** — addressed openly to the launch team. Its first sentence disclosed the method. Its operative paragraph read, verbatim:

> "Stage 1 of the #VulcanRocket will perform nominally. However, there is a risk to Stage 2 separation, trans-lunar trajectory injection and guided placement of the **Peregrine lunar lander** in the desired elliptical lunar orbit... a misguided placement... in a less than desirable orbit... will expend unplanned-for fuel/resources and jeopardize the possibility for even launching the Stage 3 of this mission."[5]

Stage 1 had performed nominally. The mission had died in the post-separation placement-and-propulsion phase, with the lander expending its propellant. The full text — including its hedges, its scrub-and-reschedule

advice, and its weather speculation, some of which had *not* happened — was preserved verbatim under a cryptographic hash, dated by two independent clocks: the platform's own timestamp (establishing **anteriority** — the post existed before the event, on infrastructure the forecaster does not control) and a Bitcoin-anchored hash of the full record manifest (establishing **integrity** — the text has not been altered since sealing).[6] Whatever this was, it was not retrodiction. The five-days-before was checkable in about a minute, by anyone, without trusting the operator.

Marsh checked it. It held.

THE RECORD, WITH ITS CEILING ATTACHED

What made the site strange to her — stranger than the method — was that its most prominent caveats were its own. A panel titled "the sample-size question, taken head-on" stated the limits she would otherwise have had to dig for:

- **n = 23.** Small by count. The record's response, in credibility-theory terms: each observation was a sealed, dated call on a named mission, several carrying nine- and ten-figure stakes, in a market whose own actuarial literature calls the entire class sparse-data. Twenty-three pre-committed launch-day observations was, the operator argued, more than the desk had from any other out-of-model source — an argument, Marsh noted, not a proof.[7] - **Self-scored.** Every grade — 22 HIT, 1 NEAR, 1 PARTIAL, 0 MISS across the 24 as of 5 August 2026, per the anchored grading ledger — had been assigned by the operator against his own pre-committed criteria. **No independent adjudication existed yet.** The re-grading harness was the operator's answer: overturn me. Nobody, as far as the site showed, had publicly tried.[3] - **The base rate ties the aggregate.** The published aggregate Brier score (0.0903 across the full 106-call corpus as of 5 August 2026, launch plus geopolitical; 0.0349 on the launch subset) was, the record's own materials conceded, roughly what a naive base-rate forecaster would score — most launches succeed, so headline calibration alone cannot separate skill from a well-chosen prior. If there was signal, it lived in the day-level *risk-class* calls — which failure mode, which phase, which window — not in the average.[7] - **The comparator was n = 0, stated as a challenge.** The operator's claim: to his knowledge, no other public, sealed, pre-launch, graded launch-day forecast record existed — from any consultancy, model vendor, or forecaster — and if one existed anywhere, name it and the page would link it. A falsifiable open challenge, not an established fact. Until falsified, the operator argued, the

choice was not "23 versus more"; it was 23 versus none.[7] - **And the rail.** Printed where the money was tabulated: *"No prevention claim is made or implied — what any read would have changed is unknowable and never asserted. The day is readable, not preventable."* The record did not claim it could have saved Peregrine, or the $30 million satellite, or anything else. It claimed only to have said what it said, before the event, on the record.[8]

The money exhibits she could verify against her own market's numbers, and they checked: the one insured loss in the set was real and the claim had been paid[2]; the delay ledger's figures traced to NASA's Inspector General and GAO[9]; the rate history was her own trade press.[1]

THE AFTERNOON IN QUESTION

The junior analyst had estimated the full back-test at **one afternoon**: download the CSV, run the verifier, spot-check five sealed texts against their hashes and platform timestamps, re-grade the three or four most consequential calls against the operator's published criteria, and write a one-page internal note. Perhaps four hours of a junior's time — call it a few hundred pounds of desk cost against a book where a single point of rate on a single placement was worth more than that by orders of magnitude.

Marsh listed, honestly, the reasons not to spend it:

Opportunity cost with a queue attached. The desk received a steady stream of models, consultancies, and cranks, each claiming an edge. If the bar for an afternoon was "the claim is checkable," the desk would do nothing but check claims. Attention was the scarcest asset she managed.

Adverse selection of weird claims. Claims that reach an underwriter's inbox with astonishing hit rates are disproportionately the survivors of some selection process — lucky streaks self-publish; unlucky ones delete their accounts. Twenty-three-for-twenty-three *sittings* was exactly what a survivorship pipeline would deliver to her desk. (Though — she made herself note — the two-clock sealing made *post-hoc* selection of texts harder than usual, and the manifest claimed to be the complete corpus, a claim the verifier could partially test. Harder was not impossible.)

The base rate of extraordinary claims. Nearly every extraordinary-accuracy claim in the history of forecasting had dissolved on inspection — into vague predictions retro-fitted to events, cherry-picked windows, or grading generosity. Her prior that *this* one dissolved too was, rationally,

overwhelming. And the record's own concession — that the aggregate Brier ties the base rate — showed where the dissolution would probably be found: in the grading of the risk-class calls, precisely the part that was self-scored.

Reputational contagion. "Lloyd's-market desk evaluates astrologer" was a headline that wrote itself. Even a private memo leaks; even a negative finding associates the desk with the subject. Institutions run on the company they are seen to keep, and there was no upside scenario in which she could book the finding as ordinary due diligence. Engaging unacceptable sources has a cost that is real even when the engagement is skeptical.

The mechanism prior. She had no causal model under which planetary positions inform launch outcomes, and adopting an input without a mechanism was how desks bought noise.

Then she listed, honestly, the reasons to spend it:

The claim was checkable, and checkable claims are rare. Almost everything in her inbox required trusting the sender. This one, uniquely, did not: the anteriority was on infrastructure the operator didn't control, the integrity was on a blockchain, the grades came with a harness for overturning them, and the whole kit was free. The Peregrine entry alone — a named payload, a named phase, five days out, on a mission her market remembered — was either the most elaborate survivorship artifact she had seen or a data point with no peer. One afternoon distinguished the two.

Her own profession's logic pointed at the test, not the prior. Experience rating — her desk — was precisely the practice of updating on demonstrated records without demanding a mechanism first. The market had repriced Falcon 9 on its record, not on a theory of Merlin engines. She did not need to believe anything about astrology to ask the only question underwriting ever asks: *is the record real, and does it discriminate?* Feynman's formulation, which she half-remembered from university, was blunter: the test of all knowledge is experiment — and a claim that volunteers its own experiment, pre-registered and adversarially verifiable, was making itself easier to kill than any claim she'd seen that quarter. Refusing to run a free experiment was not skepticism; it was something else wearing skepticism's clothes.

The asymmetry of the payoff. If the back-test dissolved the record — the likely outcome, on her priors — the afternoon bought a definitive internal answer and a template for the next weird claim. If it *didn't* dissolve — if the seals held, the grades survived re-grading, and the risk-class calls showed

discrimination beyond base rate at n=24 — then she held information almost nobody in her market had priced, about a class of risk (the day itself: scrubs, slips, anomaly windows) that her market's own actuarial literature admitted it reserved badly.[9] Small probability, large value, bounded cost: her desk priced exactly this shape of bet every week — in the other direction.

The n=0 comparator, if it held, was itself market information. If it was really true that no consultancy or model vendor anywhere had ever published a sealed, graded, pre-launch record — that the entire launch-forecasting industry, conventional and otherwise, ran unsealed — that fact alone said something about her information environment worth an afternoon.

And the disclosure cut both ways. The operator had put the method in the first sentence — the single least persuasive way to package the record. A survivorship huckster, she reasoned, buries the tell; this one led with it, published his misses framework, conceded the base-rate tie, and asked to be re-graded. None of that made the record true. It made the record *cheap to check* — which was a different property, and the one her decision actually turned on.

THE DECISION

It was 2:40 p.m. The analyst was waiting for an answer on whether to spend the rest of the day on it. The renewal book was open on her other screen. Whatever she decided would be, she understood, a policy and not an incident: whatever standard she applied to this artifact, she was implicitly adopting for every checkable-but-unacceptable claim that reached the desk after it — and for every acceptable-but-uncheckable one too.

Marsh looked at the two clocks on the Peregrine entry one more time — the platform's timestamp and the Bitcoin anchor, five days before a loss her market remembered — and picked up the phone to give the analyst an answer.

Does the desk spend the afternoon?

EXHIBIT 1 — THE SEALED PEREGRINE ENTRY (LA-003), AS PUBLISHED

Sealed: January 3, 2024 (platform-timestamped public post, ID 1742705351995322731) — five days before the January 8, 2024 launch. **Method disclosure, first sentence of the sealed text:** "I'm a Vedic Astrologer..." **Grade**

(self-scored, per the anchored ledger): HIT. **Pre-event probability, frozen at seal:** 0.85. **Record hash:** 7ea95cf75dbd5a0b25d494f7fb5b2aad4a000566c147ef66e38e588c9307566f.

Operative text, verbatim (honest ellipses; full text public at the source):

> "Stage 1 of the #VulcanRocket will perform nominally. However, there is a risk to Stage 2 separation, trans-lunar trajectory injection and guided placement of the Peregrine lunar lander in the desired elliptical lunar orbit... a misguided placement... in a less than desirable orbit... will expend unplanned-for fuel/resources and jeopardize the possibility for even launching the Stage 3 of this mission."

What happened: Vulcan's certification flight was flawless (Stage 1 nominal, as stated). Hours after separation, Peregrine suffered a propulsion anomaly, expended propellant in corrective maneuvers, lost the Moon landing, and re-entered on January 18, 2024.

Honesty note carried on purpose: the full sealed text also advised that a launch at the planned time faced impediments and suggested weather or ground-system delays — the launch in fact lifted off on schedule. The record's grade rests on the operative failure-class claim quoted above, per the operator's published grading criteria; the full unexcerpted text, including the parts that did not happen, is preserved under the hash and is exactly what the re-grading harness invites readers to re-score.

Sources: jyotishintelligence.com/advisory/LA-003 (full verbatim sealed text, hash, grade, criteria); NASA/Astrobotic post-mission review coverage via GeekWire, geekwire.com/2024/astrobotics-peregrine-moon-lander-mission-end/.

EXHIBIT 2 — THE RECORD AT A GLANCE, WITH ITS PUBLISHED CEILING

Fact	Value	Caveat published by the record itself
Graded launch-day calls	23 (2023–2026: Chandrayaan-3, Starship IFT-2/3, Vulcan Cert-1, New Shepard NS-24, Axiom-3/4, NISAR, GSAT/LVM3, Artemis II windows, New Glenn NG-3)	n is small; offered as back-test + fixed-criterion forward season, never a rate input
Grades (per anchored ledger)	22 HIT · 1 NEAR · 1 PARTIAL · 0 MISS	**Self-scored** by the operator against pre-committed criteria; **no independent adjudication yet**; re-grading harness public

Aggregate Brier (full 106-call corpus)	0.0903	**A naive base-rate forecaster roughly ties this** — the aggregate alone cannot separate skill from prior; discrimination, if any, is in the day-level risk-class calls
In-class risk flags realized	15 of 15 (operator's count) → one-sided 95% binomial lower bound ≈ 0.82 on in-class hit propensity	The operator's own arithmetic on the operator's own grades — exactly what the harness lets a reader recompute
Comparator	**n = 0** — no other public, sealed, pre-launch, graded launch-day record known	Stated by the operator as a **falsifiable open challenge** ("name one and this page will link it"), not as established fact
Prevention claims	**None**	"The day is readable, not preventable" — printed on the record's own money page

Source: sealed manifest, anchored grading ledger, and sample-size panel at jyotishintelligence.com (all figures are the record's own published figures, reproduced with their published caveats).

EXHIBIT 3 — THE MONEY ON THE GRADED MISSIONS (THE DESK'S OWN REALITY CHECK)

Abridged from the record's sourced ledger; every figure links to a public primary or trade source at the record site.

Mission	Sourced cost	Outcome	Who paid on the failure
Vulcan Cert-1 · Peregrine	$108M CLPS award	Lander lost after flawless launch	US taxpayer, ≈$97M (only 10% of award success-gated); builder absorbed its own loss on top[a]
New Glenn NG-3 · BlueBird 7	$30M insured value	Total loss — unsurvivable orbit; FAA grounded the vehicle	**An insurer — the one commercially covered asset in the set; claim paid**[b]
Artemis II · SLS/Orion	$4.1B/launch	Flew ~17 months late (beside $4.79B FY2024 obligations — two facts, never multiplied)	No failure; the taxpayer carried the slip[c]
NISAR · GSLV-F16	$1.5B	Flew >1 year late; $30M GAO-logged growth	Two governments, uninsured[d]
Starship IFT-2/3	unpublished (FAA TPL layer $500M/flight)	Both test stacks destroyed	Provider — a failed Starship test is not an insured loss (London market)[e]
Aggregate	**~$6.2B at stake under the 23 calls**	$138M destroyed on the public record	**$30M = the only insured claim**

Sources: (a) GeekWire/NASA review coverage. (b) Payload owner's Form 8-K (SEC); Space Intel Report. (c) NASA OIG IG-22-003; planetary.org budget compilation. (d) GAO-23-106021. (e) FAA license VOL 23-129; The Insurer. Full

sourced table: jyotishintelligence.com.

EXHIBIT 4 — THE TWO-CLOCK VERIFICATION ARCHITECTURE (WHAT ONE AFTERNOON BUYS)

Layer	Clock	What it establishes	How the desk checks it
Platform timestamp	The publishing platform's own record (post IDs, dates) — infrastructure the forecaster does not control	**Anteriority** — the text existed before the event	Open the post; compare its date to the launch date (~1 min per call)
Bitcoin-anchored manifest hash	OpenTimestamps attestation of the record manifest on the Bitcoin blockchain	**Integrity** — no sealed text has been altered since anchoring	Run the zero-dependency verifier script against the public manifest (~5 min)
Frozen grading ledger	Separately anchored ledger of grades + pre-event probabilities	The Brier arithmetic is reproducible, not editable	Recompute the score from the CSV (~10 min)
Re-grading harness	Reader-side	The operator's grades can be adversarially overturned	Re-score the consequential calls against the published criteria (~2–3 hrs)

The operator's division of labor between the clocks is explicit: the platform proves *before*; the blockchain proves *unchanged*. Total estimated cost of the full back-test: one junior afternoon. Price: free. Source: jyotishintelligence.com (verifier, manifest, ledger, and reproduction pages are public).

EXHIBIT 5 — HOW THE PAYING MARKET TREATS DEMONSTRATED RECORDS (THE DESK'S OWN NUMBERS)

Figure	What it shows
5.5% → 3.1%	Falcon 9 launch-cover rate, 2016 → 2019, as its record accrued (Ariane 5 eased 3.25% → 2.75%)[a]
>20% → ~7%	Rates after the 2001–03 loss cluster vs. good-record eras — the market's own cycle[b]
PSLV ↑	"The LVF rates are dependent on which rocket model is being flown and its recent reliability record" (verbatim, 2025); Falcon 9 down, PSLV up after a failure[c]

Caveats: rates also swing on market capacity, not record alone. The exhibit's point is narrow: **the market already reprices demonstrated records without demanding mechanisms** — experience rating is the desk's existing epistemology. Whether that epistemology extends to Exhibit 2 is the case's question, and this exhibit takes no side. Nothing in this case states or implies that any input lowers any premium.

Sources: (a) SpaceNews, spacenews.com/big-claims-record-low-rates-reshaping-the-space-insurance-game/. (b) Via Satellite, Nov 2007. (c) Seradata via SatNews, 17 Aug 2025.

EXHIBIT 6 — THE DECISION PACKAGE, PRICED (RECOMMENDATION LEFT TO THE STUDENT)

	A · Bin it	B · The afternoon back-test	C · Full actuarial review
What it adds	nothing	a definitive internal answer, either way	ceremony around the same answer
What it forecloses	the cheap look	nothing	weeks of desk time
Cost if the record dissolves	0	four junior hours + a template for the next weird claim	committee fatigue
Cost / value if the seals hold	information nobody in the market priced	held early, quietly	held late, loudly
Speed	—	one afternoon	a quarter
Authority required	none	her own	actuarial committee
Reversibility	n/a	full	n/a
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading)*		

CASE FOOTNOTES

1. SpaceNews (F9 5.5%→3.1%, Ariane 3.25%→2.75%); Via Satellite, Nov 2007 (>20% post-2001–03; ~7% good-record eras); Seradata via SatNews, 17 Aug 2025 (verbatim rate-follows-record quote; PSLV up after a failure).
2. BlueBird 7 (New Glenn NG-3, April 2026): $30M insured value per the payload owner's Form 8-K (SEC EDGAR); total loss and claim payment per Space Intel Report; FAA grounding per the same coverage. **No insurer is named in this case.**
3. The back-test kit — sealed-record CSV with frozen pre-event probabilities, zero-dependency hash verifier, anchored grading ledger, and re-grading harness — is public and free at jyotishintelligence.com (reproduction and evaluation pages).
4. NASA/Astrobotic post-mission review coverage, GeekWire, geekwire.com/2024/astrobotics-peregrine-moon-lander-mission-end/ — $108M award, 10% success-gated, ~$97M milestone-paid; the CEO's public statement that the mission cost the company more than it made. *Astrobotic is not criticized in this case; it absorbed a real loss of its own.*

5. Sealed advisory LA-003, jyotishintelligence.com/advisory/LA-003 — verbatim text under hash 7ea95cf7...07566f, platform post ID 1742705351995322731, dated 3 Jan 2024.
6. Verification architecture per the record's public documentation: platform timestamps establish anteriority; OpenTimestamps Bitcoin anchoring establishes integrity of the sealed texts and the grading ledger.
7. Sample-size panel, grading ledger, and comparator claim as published at jyotishintelligence.com; the 0.82 binomial lower bound and the 15-of-15 count are the operator's own published arithmetic on his own grades. The Brier figure (0.0903, n=106 full corpus as of 5 August 2026) is from the anchored grading ledger, which is append-only; recompute against the manifest and ledger hashes cited rather than the printed figure.
8. "The day is readable, not preventable" — the record's own stated rail; no prevention or counterfactual claim appears in the record or in this case.
9. NASA OIG IG-24-013 (CLPS: $208.2M increases, 14-month average delays); GAO-23-106021 (NISAR); the space-insurance trade literature's treatment of launch delay as a standing reserving problem (see the record's sourced delay ledger).

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/lloyds-market
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — Lloyd's Market: Grading an Unconventional Forecast

AUTHORSHIP DISCLOSURE. This note, like the case, was written by **Vijay Jyotish**, the operator of the forecast record the case examines. That is an unusually strong author interest, disclosed on page 1 of the case, and it should be *used* in class rather than merely acknowledged: the case hands students a live instance of its own question (see §5.5 and Assignment Question 7). All record and money facts are real and publicly checkable; the protagonist is a labeled composite; no insurer is named; the method (Vedic astrology) is disclosed plainly and is neither mocked nor advocated anywhere in the case or this note. Hard rails for the classroom: **no prevention or counterfactual claims ("the day is readable, not preventable"); no one is smeared — not the lander's builder (whose CEO publicly said the mission cost more than it made), not NASA's CLPS model (described neutrally), not any insurer, and not the forecaster; the record's limits (n=23, self-scored, no independent adjudication, base-rate ties the aggregate Brier) stay visible at all times. This note steel-mans both positions and takes no side. The question is the product.**

1. SYNOPSIS

Elena Marsh, a composite London-market space underwriter, receives a link to a real, publicly checkable artifact: a sealed, graded record of 23 launch-day forecasts (21 HIT, 1 NEAR, 1 PARTIAL, 0 MISS by its own anchored, self-scored ledger) whose disclosed input method is Vedic astrology. The record's flagship entry named "guided placement of the Peregrine lunar lander" as the at-risk phase five days before the January 2024 mission — which was lost in exactly that phase after a flawless launch that the same sealed text had also called ("Stage 1… will perform nominally"). The record is protected by a two-clock architecture (platform timestamps = anteriority; Bitcoin anchoring = integrity), publishes its own ceiling (small n, self-scored, base-rate ties the aggregate Brier, comparator n=0 stated only as a falsifiable challenge), makes no prevention claims, and offers a free back-test: a CSV, a zero-dependency verifier, and a re-grading harness — estimated cost, one junior afternoon. The case inventories Marsh's honest reasons not to check (opportunity cost, adverse selection of surviving weird claims, the base rate of extraordinary claims, reputational contagion, no mechanism) and her honest reasons to check (the claim is checkable and checkable claims are rare; experience rating — her own craft — updates on records without mechanisms; asymmetric payoff at bounded cost; the n=0 comparator is market information if true; the artifact is cheap to kill). It ends at 2:40 p.m., with the analyst waiting: **does the desk spend the afternoon?**

2. LEARNING OBJECTIVES

Students should leave able to:

1. **Separate three questions that instinct fuses into one:** Is the claim true? Is the claim *checkable*? Is checking it *worth the cost*? — and recognize that professional decisions turn on the second and third far more often than on the first.
2. **Construct a value-of-information analysis** for a bounded-cost test of a low-prior, high-payoff claim — and identify when "my prior is ~0" does and does not justify not testing (it depends on the cost of the test and the value of the tail, not on the prior alone).
3. **Articulate the adverse-selection/survivorship critique** of inbox-arriving track records, and evaluate how much of it pre-registration and two-clock sealing does and does not neutralize.
4. **Name reputational risk as a real decision input** — institutions rationally price the company they are seen to keep — and distinguish it from an epistemic reason to disbelieve.
5. **Apply the Feynman standard** ("the test of all knowledge is experiment") to a professionally uncomfortable case, and notice when "skepticism" functions as a reason to avoid an experiment rather than to run one.
6. **Audit self-scored evidence:** understand why self-grading is the record's weakest link, why the base-rate tie on the aggregate Brier concentrates the entire question into the risk-class grading, and what an adversarial re-grade would actually test.
7. **Read interested sources:** practice extracting verifiable content from a maximally interested author (including the case's own author) without either credulity or reflexive dismissal.

3. TARGET COURSES

- **Managerial Decision-Making / Judgment under Uncertainty** (flagship placement) — priors vs. evidence, value of information, noise vs. signal. - **Forecasting courses (Tetlock-adjacent syllabi)** — pairs naturally with Superforecasting material: pre-registration, Brier scoring, base-rate comparison, the difference between calibration and discrimination; the record is a live specimen of scoring-rule mechanics, whatever one concludes about it. - **Insurance / Underwriting programs** — experience rating as an epistemology; sparse-data classes; what counts as a rateable record. - **Business Ethics — epistemic due diligence** — what professionals owe checkable claims; reputational risk vs. truth-seeking; the ethics of dismissal-without-inspection and of engagement-without-endorsement. - **Research methods / evidence-based management** — pre-registration, adversarial verification, survivorship bias.

Level: MBA, EMBA, and advanced undergraduate. The case requires no space or insurance background and no position on astrology — instructors should say the latter out loud in the first five minutes.

4. ASSIGNMENT QUESTIONS

1. Before class, spend 15 minutes (not more) on the public record at jyotishintelligence.com. What did you check first, and why that? What would a full afternoon add?
2. Estimate, with explicit numbers, the expected value of the afternoon: your prior that the record survives adversarial re-grading; the value to the desk if it does; the cost if it doesn't (include reputational cost honestly). What prior would make the afternoon not worth it?
3. Steel-man the **don't check** position in one page, using only arguments you would defend in front of the class. (Weak versions — "astrology is silly" — will be returned.)
4. Steel-man the **check** position in one page, under the same rule. (Weak versions — "it's 23-for-23, it must work!" — will be returned; the aggregate Brier ties the base rate, and the record says so itself.)
5. The record's grades are self-scored. Design the afternoon: which calls do you re-grade first, against what criteria, and what result would make you (a) stop early, (b) escalate to a forward test? What does the fixed-criterion forward season the operator offers solve that the back-test cannot?
6. Marsh notes her decision is "a policy, not an incident." Write the desk's one-paragraph standing policy for checkable-but-unacceptable claims — one you would also apply to acceptable-but-uncheckable ones.
7. The case's author is the record's operator. Identify two places where that interest could have shaped the case despite the sourcing, and two facts the interest cannot have contaminated (and say why).
8. **The machine staff-officer.** An AI diligence agent can now run Marsh's entire afternoon — download the CSV, verify every hash, re-grade every call against the published criteria — in ninety seconds, for every strange claim that reaches the desk. The bounded-cost objection is gone. What replaces it as the binding constraint (attention? the weighting decision after verification?), and which step of Exhibit 4's two-clock architecture must still terminate in a human judgment? Write the desk rule.

5. ANALYSIS (INSTRUCTOR'S ROADMAP)

5.1 FRAME THE REAL QUESTION EARLY, OR LOSE THE ROOM

The session fails in two symmetrical ways: the room ridicules the method for eighty minutes, or the room argues about whether astrology is real. Both are the same failure — collapsing *is it true?* into the seat where *is checking it worth an afternoon?* should sit. Open by writing the three questions (§2, LO1) on the board and announcing that class will spend the day on questions two and three, because question one is precisely what nobody in the room can answer without the afternoon — that's what makes it a case.

5.2 THE DON'T-CHECK POSITION, AT FULL STRENGTH

Give this position genuine power; it is not the villain's side. Its strongest form:

- **Attention is the binding constraint, and the queue is adversarial.** A desk that checks everything checkable gets selected against: the population of claims arriving at inboxes is generated by processes that optimize for arriving at inboxes. Survivorship does the sender's marketing: 23 good sittings self-publish; the streaks that died did not build websites. A rational desk prices the *generating process* of its inbox, not the artifact in front of it. On that pricing, the correct per-item engagement is often zero — not because any given claim is false, but because the expected yield of the class is below the cost of triage. This is the same logic by which the desk deletes cold pitches from conventional quants with backtested "alpha." - **The base rate of extraordinary claims is brutal.** Centuries of accuracy claims from unconventional methods have dissolved on inspection — and, more pointedly, most *conventional* forecasting edges dissolve on inspection too. The record's own concession (base-rate ties the aggregate Brier) tells you where this one would dissolve: in self-scored grading of ambiguous risk-class language. The sealed texts are long and hedged; the LA-003 entry also predicted launch impediments that did not occur (the case's Exhibit 1 says so on its face). Grading generosity is exactly the degree of freedom self-scoring leaves open, and n=23 with one grader is not evidence that survives that critique — it is the critique's home turf. - **Reputational contagion is a real cost, rationally priced.** The desk's franchise is institutional trust. "Syndicate evaluates astrologer" has no good-news version in the trade press; even a debunking memo, leaked, becomes "they took it seriously enough to check." Institutions are not cowardly for pricing this — they are pricing a genuine externality on colleagues and capital that the individually curious underwriter does not bear alone. Note for the class: this is a *prudential* reason, and it is fully compatible with believing the record might survive checking. That compatibility is uncomfortable, and the discomfort is the lesson. - **No mechanism means no out-of-sample confidence even if the back-test holds.** Without a causal story, a clean 23 gives you no basis to expect call 24 — the reference class for "how long do mechanism-free streaks continue" is grim. The afternoon cannot buy what the desk actually needs, so its information value is lower than it looks; only the forward season could, and that is no longer one afternoon.

5.3 THE CHECK POSITION, AT FULL STRENGTH

- **Checkability is the rare property, and it inverts the survivorship math — partly.** Survivorship explains streaks that are *presented*; it struggles more with streaks that are *pre-registered on infrastructure the claimant doesn't control and offered for adversarial re-grading at zero price*. The two-clock architecture doesn't prove skill; it converts "trust me" into "audit me," collapsing the cost of resolution from unbounded to one afternoon. A desk that won't spend a bounded cost to resolve an audit-me claim has revealed that its skepticism is not primarily about evidence. (Honest counter, keep it on the board: sealing defeats *text* selection, not *forecaster* selection — many sealed forecasters could exist, with only the lucky one mailing underwriters. The manifest's

completeness claim and the n=0 comparator bear on this but don't close it. This exchange — seal vs. survivorship, round after round — is the analytical heart of the session.) - **Value of information dominates at these stakes.** Run Q2's arithmetic openly. Suppose the class's prior that the record survives re-grading and shows real risk-class discrimination is 2%. The cost is one junior afternoon plus (bounded, mostly internal) reputational exposure. The value in the 2% branch — a day-level input in a sparse-data class the market admits it reserves badly, on a book where single rate-points dwarf the cost — is enormous, and the 98% branch still yields a reusable triage template. Most classes discover that the afternoon survives even sub-1% priors; the interesting discussion is which *reputational* cost assumptions reverse it, because that is where the real disagreement lives. - **Experience rating is the desk's own epistemology, already mechanism-free.** The market cut Falcon 9's rate on its record, not on a combustion model; it raised PSLV's on one failure (Exhibit 5). Underwriting has never demanded mechanisms before updating on demonstrated records — it demands *authenticated* records. So the desk's actual standard, applied consistently, points at the verifier, not at the mechanism prior. A student who says "but astrology can't work" should be asked what the desk's Falcon 9 repricing "worked" by, and whether the desk knew. - **Feynman's standard, stated precisely.** "The test of all knowledge is experiment" does not oblige anyone to believe the record; it obliges a professional who *cares whether it's true* to prefer the free experiment over the untestable prior. The record has made itself unusually easy to kill: frozen probabilities, published criteria, re-grading harness, an open falsifiable comparator claim, and a stated forward season. Declining a free kill-shot at a claim you're sure is false requires explanation, and the honest explanations available (attention, reputation) are the don't-check side's *real* arguments — the class should notice that "it's obviously false" quietly exits the argument once the test is free.

5.4 THE SYNTHESIS THE NOTE REFUSES TO MAKE — AND WHAT CAN BE SAID INSTEAD

Take no side on the decision. What the instructor *can* legitimately crystallize:

1. **The two positions disagree about costs, not about astrology.** By end of session the live variables should be attention policy, survivorship modeling, and reputational pricing — not the zodiac. If that migration has happened, the case has worked.
2. **Self-scoring is the fulcrum.** Everything reduces to: would the grades survive a hostile re-grader? That is answerable only by the afternoon (or by the operator obtaining independent adjudication, which the record itself lists as its missing piece). Both sides agree on this; they disagree on whose move it is.
3. **The decision is a policy.** Marsh's closing realization is the transferable asset: whatever rule she applies here, she owns everywhere — including for comfortable sources whose records are *not* sealed. Several students will notice the sharpest implication: the desk routinely acts on conventional inputs with worse verification properties than this record. That observation is allowed to stand; it is about verification properties, not about astrology, and it makes no claim the record works.

4. **No prevention claims, ever.** If a student says the read "could have saved" Peregrine or the $30M satellite, quote the rail — "the day is readable, not preventable" — and note that the record itself never makes the claim. This discipline is also a craft lesson: the record's refusal to claim counterfactuals is part of what makes it auditable at all.

5.5 THE AUTHORSHIP LAYER (DO NOT SKIP)

The case was written by the record's operator, disclosed on page 1. In the final ten minutes, turn the case on itself: *you have just spent eighty minutes inside a decision about whether to check an interested party's checkable claims — while reading an interested party's checkable case.* Q7 operationalizes it: the money figures, the OIG/GAO findings, the insurance-rate history, the platform timestamps, and the hash are all outside the author's control and independently verifiable; the composite's inner monologue, the selection of which caveats to foreground, and the very framing "the question is the product" are not. Students who conclude "so I should verify the load-bearing facts and discount the framing" have derived the case's answer to itself — for both artifacts, the case and the record. Note the recursion honestly: the author benefits if the class decides checking is virtuous. The class's protection is the same as the desk's — the facts are checkable, and the assignment (Q1) already made them check.

5.6 COMMON TRAPS

- **Mockery** ("we're pricing horoscopes now?") → restate the no-mock/no-advocate frame; redirect to LO1's three questions. - **Credulity** ("23-for-23 — how is this even a question?") → the record's own ceiling: self-scored, n=23, aggregate Brier tied by the base rate; the hedged full texts (Exhibit 1's honesty note) are exactly where a hostile re-grade would attack. - **Prevention claims** → the rail, every time: readable, not preventable. - **Naming/smearing** → no insurer is named; the builder absorbed its own loss and is off-limits; the forecaster is not to be called a fraud in class — the record is either audited or it is not, and unaudited dismissal is precisely the behavior under examination (as is unaudited belief). - **"Just do a forward test instead"** → good instinct, incomplete: the forward season needs pre-set criteria and months; the back-test is this afternoon and shapes whether a forward test is warranted. Sequencing is part of Q5's answer.

6. BOARD PLAN — 75–90 MINUTES

Time	Segment	Board
0–5	Ground rules: no position on the method will be taken or required; no prevention claims; nobody gets smeared. The three questions (true? / checkable? / worth checking?) go up and stay up.	Board 1: the three questions
5–15	**The artifact, deadpan.** Walk Exhibit 1 (LA-003, two clocks, verbatim text, the honesty note about the parts that didn't happen) and Exhibit 2 (the ceiling). Poll #1, hands: does the desk spend the afternoon? Record the split.	Board 2: the artifact's anatomy — and its published limits

15–35	**Steel-man DON'T CHECK** (Q3): harvest and strengthen — attention policy, adverse selection of inbox claims, extraordinary-claims base rate, reputational contagion, mechanism-free out-of-sample problem. Push until the room agrees the position is respectable.	Board 3, left column
35–55	**Steel-man CHECK** (Q4): checkability as the rare property; the seal-vs-survivorship exchange (run at least two rounds); the VoI arithmetic from Q2 with the class's own numbers; experience rating as the desk's existing mechanism-free epistemology; Feynman's free experiment.	Board 3, right column; Board 4: the VoI calculation
55–65	**The fulcrum:** self-scoring. Design the afternoon together (Q5): which calls first (LA-003 and the highest-stakes windows), what criteria, stop/escalate rules, and what only a forward season can establish.	Board 5: the re-grade protocol
65–75	**Policy, not incident** (Q6): draft the desk's standing rule; test it against a comfortable-but-unsealed conventional vendor. Poll #2: same question. Compare with Poll #1 and ask movers *what moved them* — costs, never the zodiac, if the session worked.	Board 1 annotated: the migration of the argument
75–90	**The authorship turn** (§5.5, Q7): the case is its own specimen. Which facts could the author's interest not contaminate? What did you verify? Close by refusing the verdict explicitly: the note takes no side; the question is the product. Last line on the board: *readable, not preventable — and checkable, which is a different property from true.*	Corner board

For a 90-minute executive session, expand the VoI segment with real desk-cost numbers from the room and compress Q6.

7. EPILOGUE (RELEASE AFTER DISCUSSION)

As of mid-2026, everything the case describes remains publicly standing at jyotishintelligence.com: the sealed manifest and Bitcoin anchors verify; the grading ledger holds 23 launch-day calls (21 HIT, 1 NEAR, 1 PARTIAL, 0 MISS, self-scored, frozen probabilities); the back-test kit remains free; the comparator claim — no other public, sealed, pre-launch, graded launch-day record known — has not been publicly falsified; and the operator's stated offer of a fixed-criterion forward season stands. **No underwriting desk, syndicate, or institution is publicly known to have run the back-test.** No independent adjudication of the grades exists yet — the record's own listed missing piece. The record has since continued to seal forward calls, each of which will grade itself in public whether or not anyone is checking.

The one insured loss in the graded set — the $30M satellite — was paid by its (unnamed) insurer. Astrobotic continued flying CLPS missions. Elena Marsh does not exist, and what her desk decided at 2:40 p.m. is therefore, permanently, the reader's problem.

CASE THREE

The Xi Weekend: A Counterparty Question

A teaching case in political risk, supply-chain exposure, and the discipline of the documented hold

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author therefore has a direct interest in that record. Every date, post, and public event in this case is real and independently checkable at the sources in the exhibits. The **protagonist is a composite character** — "Maya Chen" is not a real person; her firm, meetings, and deliberations are a teaching device, and no fact she reviews is invented. The case is prepared as a basis for class discussion and takes **no position** on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly and publicly: **Vedic astrology.** This case neither mocks nor advocates the method; it treats the record as a document with dates and grades beside it, and leaves the epistemic question to the classroom.

SATURDAY, SEPTEMBER 24, 2022 — 07:40, CUPERTINO

Maya Chen's phone had been buzzing since six. As vice president for supply-chain risk at a mid-cap consumer-electronics firm with roughly seventy percent of its assembly base in mainland China, she was the person a certain kind of weekend found. This was that kind of weekend.

The rumor had started moving on Friday: *Xi Jinping under house arrest in Beijing.* By Saturday morning it was not a whisper but a wave. A former Indian cabinet minister with more than five million followers had posed it to the world as a question worth checking.[1] An Indian television channel was running it as breaking news — *#BREAKING: Xi Jinping under house arrest… Beijing under army control* — with a named successor, a general, attached.[2] The rumor's supporting cast had assembled itself overnight: claims of mass flight cancellations, troop columns near the capital, a purged security establishment. In eighteen days, the 20th Party Congress would convene to

decide whether Xi received a precedent-breaking third term.[3] If there were ever a plausible week for the unthinkable in Beijing, the world's amateur Kremlinologists agreed, it was this one.

Chen's problem was not deciding whether the rumor was true. Her problem was that *no one could* — not on Saturday. The fact-checking organizations had published nothing yet.[4] The embassies were silent. Beijing itself was silent, which the rumor's proponents read as confirmation and its skeptics read as a weekend. Every channel Chen would normally use to resolve a question like this either had nothing or was repeating the rumor back to her.

WHAT THE DESK HAD

At 09:15 her deputy forwarded her something stranger than the rumor: a public post, timestamped that morning at the panic's peak, from a source the deputy described, carefully, as "an unconventional forecaster with a public track record." The post quote-replied the five-million-follower amplification directly and stated flatly that the house-arrest and removal-from-power claims were false — not *unverified*, not *developing* — false.[5] The source's record, the deputy noted, was unusual in a way that cut through the usual anonymous-account problem: its forecasts were published dated, sealed cryptographically before events, and graded afterward in public, misses included, on a record anyone could check.[6] Its method, disclosed in its own materials, was Vedic astrology.

"So we have," Chen said, "a viral coup rumor amplified by a former cabinet minister, zero official information, and a flat *false* from an astrologer with a scoreboard."

"That's the desk," the deputy said. "What do we do with it?"

THE LADDER

By 10:00 Chen had sketched the realistic menu on a legal pad. She labeled the rungs by cost.

Rung zero — do nothing, log nothing. The firm had no China assessment on file for this weekend; if the rumor evaporated, no one would ever ask. If it didn't, everyone would.

Rung one — the documented hold: a dated log entry stating what was assessed, what was judged, and what would trigger a revisit. Cost: fifteen minutes. Fully reversible.

Rung two — convene the crisis team Saturday afternoon. Cost: a burned weekend for nine people, and the subtler cost Chen had learned to price — every convening spent a little of the team's alarm credibility, and the team had been convened twice already that year for events that fizzled.

Rung three — operational hedges: activate the alternative-sourcing review, price expedited freight for two product lines, alert the Vietnam contract manufacturer that a capacity query might be coming. Cost: real money, some of it unrecoverable, and a signal to partners that the firm was moving.

Rung four — financial hedges through the treasury desk, bought into a weekend panic at panic prices.

Rung five — call the audit-committee chair at home on a Saturday. There was no rung six, because rung five could not be climbed down from: a board told that China might be in a coup would not un-hear it Tuesday.

Against the ladder she wrote the two facts she actually possessed: a rumor that could not be verified, and a dated, falsifiable, publicly-checkable *false* from a source whose method she had no framework for evaluating — but whose incentives she could reason about. The source had bet a public, graded record on this call, in writing, at the peak. If Xi emerged deposed, the record would carry the miss forever; its own published rules said so.[6]

At 10:40 her phone lit again. It was the CFO: *"Board member saw the China thing on Twitter. Wants to know what we're doing about it. What do I tell him Monday?"*

Chen looked at the ladder, and at the two facts, and at the empty log entry blinking on her screen.

EXHIBITS

Exhibit 1 — The rumor's arc (all times US Pacific). Fri Sep 23: the house-arrest question posted by a former Indian cabinet minister (5.7M followers); Sat Sep 24: broadcast pickup as #BREAKING with a named successor; Sat–Sun: peak circulation, zero official comment; Mon Sep 26: first major fact-check published (Snopes: no evidence); Tue Sep 27: Xi Jinping appears publicly at a Beijing exhibition, first appearance since returning from the SCO summit; Sep 28: further fact-checks (PolitiFact); Oct 16–23: the 20th Party Congress convenes and confirms Xi's third term. Sources: x.com/Swamy39/status/1573544718507868160; x.com/thefirstindia/status/1573655421516386304; snopes.com/fact-check/xi-jinping-house-arrest-coup/; politifact.com (Sep 28, 2022); cnn.com (Sep 28 and Oct 22–23, 2022).

Exhibit 2 — The sealed input. A public post of September 24, 2022, quote-replying the amplification in Exhibit 1 and stating that the house-arrest and removal-from-power claims were false. Preserved at x.com/vijayjyotish/status/1573879965673586688. The source's record publishes each claim's confidence on a fixed linguistic ladder (this one maps to 0.95) and — relevant to any weighting decision — publishes its own *blind base-rate floor* for this claim class: 0.70 that "no coup" is true on priors alone, meaning the record itself prices the call's information content as modest. Record documentation: jyotishintelligence.com/intel/advisory/IA-CN-2022.

Exhibit 3 — The ladder, priced. Rung 0: $0 now; unbounded credibility cost if wrong kind of weekend. Rung 1: ~15 minutes; $0; reversible. Rung 2: ~9 person-weekends; alarm-credibility draw. Rung 3: low-six-figures partially unrecoverable; partner signaling. Rung 4: panic-priced premium. Rung 5: irreversible escalation. (Composite figures; teaching device.)

Exhibit 4 — What a "graded record" means here. The source's ledger, as of its August 2, 2026 snapshot (integrity hash 542e298e…): 106 graded calls — 82 hits, 10 nears, 5 partials, 9 misses — mean Brier 0.0903, misses retained permanently, all recomputable from published data and open-source verifiers. jyotishintelligence.com/grading-ledger.json.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Rung 0 (nothing, unlogged)	B · Rung 1 (documented hold)	C · Rungs 3–5 (activate)
What it adds	nothing	dated reason, revisit trigger, named falsifier	hedges at panic prices
What it forecloses	the artifact	nothing	alarm credibility; rung 5 can't be climbed down
Cost if the rumor dies	0	fifteen minutes	six figures+, partly unrecoverable
Cost / value if the coup is real	unbounded credibility cost	the trigger justifies Monday's climb	positioned — at panic prices
Speed	—	same morning	days
Authority required	none	her own	CFO / board
Reversibility	n/a	full	poor above rung 3
Recommendation & rationale		*(yours)*	

In hindsight	*(filled at grading)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS

jyotishintelligence.com/cases/xi-weekend

hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Xi Weekend

Rails that must survive discussion: the case takes no position on the method; the base-rate concession in Exhibit 2 is load-bearing and must reach the board; no claim that any real firm did or should have done anything on this weekend.

LEARNING OBJECTIVES

Students should leave able to (1) distinguish a documented hold from silent inaction by its three artifacts — dated reason, revisit trigger, named falsifier; (2) price an escalation ladder in both possible worlds before either is known; (3) separate an input's *probability contribution* (here: ~half a bit, by the record's own floor) from its *commitment value* (a dated, falsifiable, checkable stake versus everyone else's unaccountable "probably fine"); (4) reason about a source's incentives when its method cannot be evaluated.

ASSIGNMENT QUESTIONS

1. It is Saturday, 10:40. Write Chen's log entry — or defend writing nothing.
2. Price rungs 0 and 3 in both worlds (rumor true / rumor false). Which rung has the worst worst-case? The worst *expected* case?
3. The record's own floor says "no coup" was 70% likely on priors. Does that make the sealed call worthless? What exactly did the 0.95 add that the 0.70 didn't already contain?
4. What should Chen tell the CFO to tell the board member — and what artifact makes that answer safe to give?
5. Design the intake rule: what class of unverifiable input should reach Chen's desk at all, and what should its arrival trigger?
6. **The machine staff-officer.** An AI monitoring desk now auto-drafts a rung-1 log entry — reason, revisit trigger, falsifier — for every viral geopolitical rumor touching the firm's exposure map, hundreds per year. What breaks first (trigger quality? the meaning of a falsifier nobody chose?), and what does Chen's signature add to a machine-drafted hold? Design the human-in-the-loop rule that keeps rung 1 an act of judgment rather than a log rotation.

THE BOARD PLAN (75 MIN)

Open on the constraint (10): what can this desk actually do this weekend, alone? Build the ladder (20): students supply rungs; instructor prices the credibility currency explicitly. Introduce Exhibit 2 late (15) — after the room has committed to postures without it — and ask who changes rungs and why; surface the incentive analysis (the source bet a graded record). Force the commitment (20): every student writes the log entry or defends the blank. Close (10) on the transfer question and Q5.

WHAT HAPPENED, AND THE GRADE

Xi Jinping appeared publicly on Tuesday, September 27 — three days after the seal — and the Party Congress confirmed his third term twenty-nine days later. The record grades the call a HIT at 0.95, *and* publishes the concession that matters for teaching: against its own 0.70 blind floor, the information yield was roughly half a bit. The record retains nine misses at full weight elsewhere on its ledger; students who ask "but what if it had been wrong?" should be shown that the answer exists and is public. The posture the input supported — the documented hold, rung one — was also the cheapest rung on the board, and it was correct. The deepest closing point: it would *still have been the right purchase* if the coup had been real, because the log entry's revisit trigger was the tripwire that would have justified climbing the ladder Monday with the firm's credibility intact.

EPILOGUE

Fact-checks landed Monday and Wednesday — after the weekend in which the decision had to live. Every institution that activated on rungs three through five that weekend paid for insurance against an event that did not occur, at prices set by panic; almost none of them logged anything they could show afterward. The cheapest artifact on the board — four sentences with a date — turned out to be the only one with lasting value.

CASE FOUR

Thirty-Six Hours in Moscow: The Evacuation Ladder

A teaching case in crisis operations, continuity planning, and holding under maximal stress

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record. Every public event, date, and reported movement in this case is real and cited. The **protagonist is a composite character** — "Elena Marsh" is not a real person; her embassy, its nationality, its meetings, and its internal procedures are teaching devices assembled from publicly documented diplomatic practice, and no fact she reviews about the outside world is invented. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works, and depicts no real government's actual deliberations.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

SATURDAY, JUNE 24, 2023 — 06:10, MOSCOW

Elena Marsh had held the continuity-of-operations portfolio at a Western embassy in Moscow for two years, which meant she owned a document everyone hoped would stay theoretical: the Emergency Action Plan, with its graduated ladder — *authorized departure* of non-essential staff and families; *ordered departure*; and, at the top, the destruction of sensitive holdings that marks the point where an embassy stops planning to stay.

At 06:10 the theoretical period ended. Overnight, the Wagner Group — tens of thousands of armed men under Yevgeny Prigozhin, until yesterday the state's most effective private army — had seized the southern military hub of Rostov-on-Don, including the headquarters that ran Russia's war in Ukraine, and columns were now moving north on the M-4 highway.[1] Moscow was declaring counter-terrorism measures; there were reports of security forces digging in on the capital's approaches; the president was to address the nation.[2] Nobody — no ministry contact, no allied embassy, no analyst on any feed Marsh could

see — could say whether this was a mutiny, a coup, a negotiation by armored column, or the first hours of a civil war in a state holding the world's largest nuclear arsenal.

By 08:00 she was in the ambassador's office with the EAP binder. The ambassador's question was the ladder's question: *"Do we move?"*

WHAT THE DESK HAD

Marsh's section kept an intake file for long-lead, non-traditional warnings — a practice she had defended at budget time on the grounds that it cost twenty minutes a week. In it, flagged from March, was a public, cryptographically sealed forecast dated **March 6, 2023**, stating that a coup- or mutiny-class challenge to Vladimir Putin's authority lay ahead.[3] It named no actor, no city, no date — a fact the file's own intake note recorded bluntly at the time ("class only; unactionably broad as tasking; retain for context"). The source's record was public, graded, and carried its misses openly; its method, disclosed in its own first lines, was Vedic astrology.[4]

One hundred and ten days later, the class had arrived. And Marsh noticed what the file did to the room: the difference between *"this is unthinkable"* and *"this is the scenario we logged in March."* The first mindset improvises. The second executes.

What her section had actually done in March, at the cost of one staff afternoon, was think through the tripwire architecture for exactly this class: the written thresholds that would convert watchfulness into movement without requiring a shaken committee to deliberate at 2 a.m. *If armed columns cross a named line north of Voronezh. If the capital's airports close. If communications degrade in specified ways. If official movement controls touch diplomatic traffic.* Each tripwire mapped to a rung.

THE LADDER, UNDER REAL TANKS

The options in front of the ambassador at 08:30:

Hold with tripwires armed — full personnel accountability now, destruction plans reviewed now, movement thresholds re-confirmed in writing now; but no rung climbed until a tripwire tripped. Cost: a terrifying day of visible calm. Risk: if the situation collapsed faster than the tripwires could register, the embassy would be moving late.

Authorized departure now — request approval to move non-essential staff and families to a third country. Cost: visible to the host state and every ally within hours; reversible only over weeks; and it converts the embassy's assessment into a public signal in the middle of an information vacuum — a signal other embassies, markets, and the host government would all read.

Ordered departure now — the full first-stage evacuation, defensible if this was a civil war's first morning, nearly indefensible in retrospect if it wasn't.

The columns were reported moving fast — several hundred kilometers covered by midday, air assets engaged along the route.[1] At 14:00, word came that the Kremlin had been in contact with intermediaries; at 17:00, rumors of a deal; nothing confirmed.[5] The ambassador wanted a written recommendation by 21:00 — early evening in the capitals that would have to approve any movement — either *hold, with these five tripwires armed and these two rungs pre-cleared for instant execution,* or *move now, this rung, on this rationale.*

At 19:40 Marsh sat down with the March file, the tripwire sheet, and a blank recommendation memo.

EXHIBITS

Exhibit 1 — The public timeline, June 23–24, 2023 (Moscow time). Late June 23: Wagner leadership accuses the defense ministry of striking its camps; announces a "march for justice." Early June 24: Rostov-on-Don's military facilities under Wagner control; columns north on the M-4; counter-terrorism regime declared in Moscow; presidential address denounces "treason"; reported engagements along the route; by evening, columns halt ~200 km from Moscow following mediation; Wagner announces turnaround; stand-down by night. Sources: contemporaneous wire coverage (Reuters, BBC, AP, June 23–25, 2023).

Exhibit 2 — The sealed input. Public, cryptographically sealed forecast of March 6, 2023, stating that a coup/mutiny-class challenge to Putin's authority lay ahead; no actor, date, or location named. Sealed artifact links, grading, and the record's own intake-relevant concessions at jyotishintelligence.com/intel/advisory/IA-RU-005. Record ledger and verification: Exhibit 4 of the preceding case applies identically.

Exhibit 3 — The tripwire sheet (March draft, composite teaching device). Five written thresholds mapping observations to rungs, each with a named owner and an unambiguous trigger; drafted at a cost of one staff afternoon in the

week after the March file was logged.

Exhibit 4 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · Hold, tripwires armed	C · Authorized departure now
What it adds	nothing	pre-committed conditional movement	motion, visibly
What it forecloses	the rehearsed option	nothing until a wire trips	quiet reversal — weeks to undo
Cost if the mutiny dissolves	0	one terrible day of visible calm	a public signal read by every ally and the host
Cost / value if it is civil war	improvising at 2 a.m.	rungs pre-cleared, executed on trigger	early — and possibly right
Speed	—	thresholds re-confirmed in hours	capital-approval cycle
Authority required	none	ambassador's own	capital
Reversibility	n/a	full	poor
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/moscow-36-hours
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — Thirty-Six Hours in Moscow

Rails: no real embassy's decisions are depicted or graded; the sealed input named a class only — the teaching note must not inflate it into a prediction of Prigozhin, June, or the M-4; the value under examination is *pre-positioned thinking*, not foresight.

LEARNING OBJECTIVES

(1) The hold-with-tripwires-armed as a distinct posture — neither inaction nor movement, but pre-committed conditional movement; (2) why written thresholds beat committee judgment under stress (the 2 a.m. problem); (3) how a long-lead, low-specificity warning creates value *months before* its event — by relocating thinking from the worst possible moment to the best — even when it is useless as tasking; (4) signaling costs: why an embassy's rung two is never private.

ASSIGNMENT QUESTIONS

1. Write Marsh's 21:00 recommendation. One page. Commit.
2. The March input named no actor, date, or place. Its own intake note called it unactionable. Defend the twenty minutes spent logging it — or attack the whole intake file as theater.
3. Which of the five tripwires is hardest to specify without judgment re-entering? Rewrite it until a night-duty officer could execute it alone.
4. Authorized departure is "reversible." Price its actual reversal — diplomatically, personally, institutionally.
5. It is now March again, and you own the file. What does the *next* class-only warning trigger?
6. **The machine staff-officer.** An AI watch officer now computes tripwire states continuously from open sources — column positions, airport status, comms degradation — and pages the duty officer on any threshold. Which of the five tripwires are safe to delegate entirely, which produce false trips a machine can't contextualize, and what new 2 a.m. problem does automation create (the alert nobody believes)? Write the night order governing machine-tripped wires.

THE BOARD PLAN (75–90 MIN)

Constraint first (10): what the embassy can do without capital approval, and how long approvals take on a summer Saturday. The ladder (20), with the signaling costs drawn out — students consistently underprice the publicity of rung two. The March file (20):

introduce Exhibit 2 and its own "unactionably broad" intake note together; the debate over Q2 is the session's heart. The 21:00 memo (25): written commitment, then pairs defend opposite postures. Close (10) on the 2 a.m. problem and Q5.

WHAT HAPPENED, AND THE GRADE

By late evening on June 24 the mutiny had dissolved: mediation, a stand-down, columns reversing south — thirty-six hours, start to finish. An embassy that held with tripwires armed spent one terrible day and nothing else; every tripwire sheet went back in its binder untriggered. An embassy that had moved to ordered departure would have spent months walking it back. The record grades the March call a HIT on its stated class terms — a coup/mutiny-class challenge occurred — and its file carries the concession this note has already enforced: the claim's breadth. That pairing — graded hit, conceded breadth — is exactly what students should learn to price: not "the astrologer predicted the Wagner mutiny," which the record itself does not claim, but "a desk that logged the class in March met June 24 with its thinking already done."

EPILOGUE

Prigozhin's aircraft crashed two months later, on August 23, 2023 — a coda students will raise; keep the case's discipline and decline to build on it. The teaching point stays where the ladder was: the cheapest thing an institution bought that spring was an afternoon of thinking it almost certainly would never need. It needed it 110 days later, at eight in the morning, with the columns moving.

CASE FIVE

The Tariff Shock: Holding the Policy Weights

A teaching case in investment governance, fiduciary process, and the discipline of not acting

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record. The market events, dates, and public policy actions in this case are real and cited. The **protagonist and her institution are composite characters** — "Ruth Okafor" and the "Meridian Endowment" are not real; their portfolio, policy bands, and committee deliberations are teaching devices, and no outside-world fact they discuss is invented. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works, and nothing in it is investment advice.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

MONDAY, APRIL 7, 2025 — 07:30, CHICAGO

Ruth Okafor had chaired the investment committee of the Meridian Endowment — $2.1 billion supporting a research university — through a pandemic crash and a rate shock, and she had learned that committees do not panic the way people do. They panic *procedurally*: someone moves to convene, someone circulates a memo titled "risk considerations," and by the third emergency session the institution has talked itself into selling at the bottom with full documentation.

The emergency session on her calendar this morning had been requested Friday, after the worst two-day equity drawdown since 2020. The new tariff regime announced on April 2 had landed on markets like weather[1] — indiscriminate, global, and fast: US equities down double digits from their February highs within days, volatility at levels seen a handful of times in two decades, and every financial channel running the same chyron arithmetic about trillions erased.[2] Over the weekend, two committee members had emailed versions of the same motion: *reduce policy equity weights by up to 15 percent, raise cash, revisit when the environment stabilizes.* The endowment's

investment-policy statement set the equity band at 60 percent ±7; the drawdown had the portfolio drifting toward the band's bottom edge on market movement alone.

Okafor's problem was not whether the tariff shock was serious — it visibly was. Her problem was the oldest one in fiduciary governance: the policy bands existed precisely because a previous committee, in calm weather, had anticipated a morning like this one and pre-committed to how the institution should behave in it. The motion on the table proposed to overrule that calm-weather self at the exact moment the bands were doing their job.

WHAT THE DESK HAD

In the packet her CIO circulated Sunday night, between the custodian's drawdown report and the liquidity schedule, was an unusual attachment flagged by the endowment's risk analyst: a public, cryptographically sealed read, dated **April 2, 2025** — the day of the announcement — from a graded public forecast record, stating in terms preserved at the source that the shock would pass and the panicked repricing, not the tariffs, was the trap; the posture it supported was holding.[3] The analyst's cover note was careful: the source's method, disclosed in its own materials, was Vedic astrology; its record was published, dated, and graded with misses retained; and the analyst offered it "not as a forecast we adopt, but as the only dated, falsifiable, contrarian read in the packet — everything else in your inbox this weekend is momentum commentary that will deny having existed by July."

One committee member had already replied-all: *"Are we seriously reading astrology in a fiduciary process?"* The analyst had answered with a single line that Okafor kept rereading: *"You are reading a dated document that will be graded in public either way. The rest of the packet won't be."*

THE MEETING

At 08:00 the committee convened by video, eleven members, counsel present. The motion's sponsor led with the argument that always sounds strongest in a drawdown: *this time the mechanism is structural — a policy regime, not a sentiment cycle; the bands were built for volatility, not for regime change.* The rejoinder came from the committee's oldest member: *every drawdown re-describes itself as regime change; that is what being in one feels like from inside.*

Okafor had counsel walk the committee through what each path actually was, procedurally. **Hold the bands**: no action; document the review; rebalance *into* equities if drift breached the band floor — which is what the IPS, read plainly,

already instructed. **The motion**: a deliberate, documented deviation from policy, permitted under the IPS's exception clause with a two-thirds vote, a written rationale, and a scheduled re-entry review — a review that, she noted from the pandemic minutes, the committee had never once conducted on schedule.

The vote was called for 09:00. Okafor had ten minutes to decide what the chair would recommend, and she was conscious of writing for two audiences: the room this morning, and whoever would read these minutes in five years, knowing how it all turned out.

EXHIBITS

Exhibit 1 — The public market record, April 2–7, 2025. The April 2 tariff announcement and the immediate global market reaction: consecutive historic-scale down sessions April 3–4; volatility index prints among the highest in twenty years; broad indices in or near bear-market territory intraday by April 7. Sources: contemporaneous coverage (WSJ, Reuters, Bloomberg, April 3–8, 2025); public index data.

Exhibit 2 — The sealed input. Public, cryptographically sealed read dated April 2, 2025: the shock transient, the panicked repricing the trap, the posture supported: hold. Sealed artifact, confidence mapping, and after-the-fact grading at jyotishintelligence.com/intel/advisory/IA-MKT-002. Record ledger and verification tooling as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The IPS mechanics (composite). Equity band 60% ±7; drift rebalancing mandatory at band edges; deviation clause: two-thirds vote + written rationale + scheduled re-entry review. Pandemic-era minutes note: the 2020 deviation's re-entry review occurred 14 months late, after the re-entry price had doubled from the exit.

Exhibit 4 — The Decision Package, priced (recommendation left to the student).

	A · Hold, silently	B · Hold, with the dated minute	C · De-risk 15%
What it adds	nothing on paper	reviewed, reasoned, triggered — governance	relief, at the low
What it forecloses	the lookback defense	nothing	the mechanical recovery capture

Cost if markets keep falling	indistinguishable from freezing	the trigger fires and justifies action	vindicated — this once
Cost / value if the shock passes	luck, unprovable	governance, provable	loss realized + the re-entry the committee historically fails
Speed	—	one paragraph at the meeting	two-thirds vote + execution
Authority required	none	the minute itself	IPS exception clause
Reversibility	n/a	full	the re-entry problem
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/tariff-shock
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Tariff Shock

Rails: not investment advice; no real institution depicted; the case must not resolve into "astrology called the bottom" — the load-bearing artifacts are the IPS and the dated minute, and the sealed input's classroom function is to expose what *dated and falsifiable* means when every other input is neither.

LEARNING OBJECTIVES

(1) Policy bands as pre-commitment devices — the calm-weather self binding the panicking self; (2) the fiduciary hold as an *active, documentable decision* (the dated minute with a rationale and revisit trigger) versus freezing; (3) "this time it's structural" as the universal solvent of pre-commitment, and how a committee tests that claim without adopting hindsight; (4) input hygiene: what a dated, graded, contrarian read adds to a packet of undated consensus commentary — independent of its source's method.

ASSIGNMENT QUESTIONS

1. Write the chair's 09:00 recommendation and the minute that survives a five-year lookback in either world.
2. The sponsor says the bands were built for volatility, not regime change. Construct the strongest version of that argument — then state what evidence, available by April 7, would distinguish the two.
3. Does the sealed input belong in a fiduciary packet at all? Defend the analyst's one-line reply or write the rule that excludes it — and apply your rule to the rest of the packet honestly.
4. Price the motion in both worlds, including the re-entry problem (Exhibit 3's fourteen-month footnote is real committee behavior).
5. Draft the standing amendment: what should the IPS say about *emergency sessions themselves*?
6. **The machine staff-officer.** The endowment's new AI risk system — trained on decades of drawdown playbooks — recommends de-risking at the panic's peak, with high confidence and a beautifully argued memo. What does this case teach about machines trained on the room's reflexes, which IPS clauses must bind the model exactly as they bind the committee, and which cells of the minute can never be machine-filled? Draft the two-sentence amendment adding AI recommendations to the IPS's governance.

THE BOARD PLAN (75 MIN)

Open on Exhibit 3, not the crash (10) — the students must own the pre-commitment before feeling the panic. Then the drawdown and the motion (15). The both-worlds pricing of hold vs. deviate (20), with the re-entry failure base rate given real weight. Introduce Exhibit 2 (15) and run Q3 — expect the room to split on source versus properties; push each side to state a *general* intake rule and live with its consequences. Vote, with written minutes (10). Close (5): the two-audiences point — every minute is written for the room and for the lookback.

WHAT HAPPENED, AND THE GRADE

The episode resolved with unusual speed: within weeks the policy posture softened and markets retraced; by **June 27, 2025** — eighty-six days after the seal — the drawdown had round-tripped, and the record graded its April 2 read a HIT. A committee that held its bands captured the recovery mechanically; one that de-risked 15 percent on April 7 realized the loss near the low and then faced the re-entry decision its own history (Exhibit 3) says it would have failed. The record's file concedes what the teaching note should also enforce: "the shock will pass" is the higher-base-rate side of most drawdowns, and a single hit proves little — which is why the case's true subject was never the forecast. It was the machinery that made *not acting* a decision the institution could defend: the bands, the minute, and the trigger.

EPILOGUE

The most durable artifact of the morning, in the world where the committee held, is one paragraph of minutes: *reviewed the drawdown against policy; heard the deviation motion; noted the sole dated contrarian read in the packet and its public grading commitment; held; rebalancing per policy; revisit at quarter-end or a further 10 percent decline.* Four lines. They cost nothing, and in five years they are the difference between an institution that was lucky and an institution that was governed.

CASE SIX

The Interceptor Window: Allocation Under Scarcity

A teaching case in defense resource allocation, warning intake, and the reweight posture

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record. Every strike, date, capacity figure, and public statement in this case is real and cited. The **protagonist and his cell are composite characters** — "Major Tomas Berg" and the "Coalition Air Defence Coordination Cell" are not real; no actual military staff's deliberations are depicted, and no fact Berg reviews about the outside world is invented. The case asserts **no claim that any strike could have been prevented**, takes no position on whether the forecast method referenced in Exhibit 2 works, and grades no real institution.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

FRIDAY, MARCH 29, 2024 — 21:15, SOMEWHERE IN EUROPE

Major Tomas Berg ran the intake-and-priorities desk of a small multinational cell that helped coordinate air-defence support to Ukraine — resupply sequencing, battery rotation advice, and the unglamorous spreadsheet at the center of the war's hardest arithmetic: the defended-asset list. The list was always longer than the coverage. That was not a planning failure; it was the definition of the job. Integrated air and missile defence under scarcity is dispatch — the same problem a grid operator solves with megawatts, solved instead with interceptors, where being short at the wrong node for one night is not a brownout but a crater.

Tonight the arithmetic was worse than usual. Russia had hit Ukrainian energy infrastructure with major waves on March 22 and again this morning — generation targets, not transmission, a shift the analysts had flagged.[1] Interceptor stocks were publicly, painfully low; the debate over resupply was playing out in parliaments, not in his cell. And on Berg's screen was the April priorities memo he owed the cell's director by Monday: the recommended

defended-asset weighting for the coming month, and whether to convene the contingency session that allied energy-security staffs had penciled but never scheduled.

Under the memo, flagged by his own six-month-old note, was the strangest item in his intake file.

WHAT THE DESK HAD

Berg's cell, like every warning function, kept a lane for inputs it could not grade: tips, walk-ins, model outputs, and — since a junior analyst had argued it cost nothing to log — dated public forecasts from a source with a sealed, graded record. The item in question was a public post of **September 4, 2023**, item T32 in a numbered sequence, cryptographically timestamped, which read in the words preserved at the source: *"There are likely to be massively violent escalations during March 9-11, 2024 and April 10-12, 2024. The gloves will be off at this time as the final knock-out blows are exchanged in this conflict."*[2]

Berg's own intake note, written in September, had been disciplined about what the text did and did not contain: *"Two three-day windows, named 187 and 220 days out. Violence-intensity claim in battlefield language. Names Russia. Names NO target, NO technique, NO axis. Unactionable as tasking. Retain; revisit if window approaches with campaign context."* The source's method, disclosed in its own materials, was Vedic astrology; its record was public, graded, and kept its misses at full weight.[3]

The window was now twelve days away. And the campaign context had arrived on its own: two mass waves against generation in the past week. The March 9–11 window in the same post had, for what it was worth, coincided with a violent stretch — the Odesa strike that landed near a presidential motorcade among them.[4]

Berg was rigorous with himself about the seam, because the seam was where careers and analysis both went to die: **the seal supplied a calendar. It did not supply a target.** "Kyiv-region thermal generation" was not in the September text; it was Berg's own threat-model inference from an eighteen-month public campaign and two fresh waves. If he wrote the memo weighting thermal generation for April 10–12, the hard half of that judgment was his, and he would own it either way.

THE THREE PACKAGES

He drafted the options the way his cell's doctrine required — priced, bounded, and honest about authority.

Package A — Monitor. The window enters the daily brief as context; no posture changes. Cost: nothing now. Price if the window is real: the cell spent its warning on a sentence.

Package B — The bounded reweight. Recommend that for the 72 hours of April 10–12, allocation guidance weight interceptor coverage toward the asset class the campaign was already pointing at — Kyiv-region thermal generation — accepting thinner coverage elsewhere for exactly that window; convene the penciled contingency session in the coming week, so resupply sequencing and battery-rotation advice could be aligned before, not after. A tasking recommendation inside existing authority. Cost: one movement cycle, 72 hours of redistributed risk, and the credibility spent recommending it on this evidentiary basis. Reversible by design: the dated window is what caps the cost at 72 hours instead of "indefinitely."

Package C — The elevated posture. Push for a standing reweight toward generation assets for the whole spring. Cost: exactly the allocation error a saturation adversary wants to induce — you cannot defend everything always, and an indefinite posture converts a bounded hedge into an open-ended bet against every other asset on the list.

Against Package B stood the objection Berg wrote into his own draft because it was true: Ukraine's binding constraint was **absolute scarcity**, not allocation timing. A reweight moves scarcity around; it does not manufacture interceptors. Against an 80-munition saturation package, a redistributed defence might simply fail at a different address. And a further strike on generation was close to a base-rate expectation by late March — everyone knew more was coming. What nobody else had put in writing was *when*. The only dated, falsifiable claim in his entire intake stream — the only input whose lead time was not capped by the adversary's own observable preparation clock — was the strange one, and the part of the problem it addressed, the calendar, was precisely the part an allocation planner can act on 72 hours at a time.

The memo was due Monday. Berg opened it and typed the header: *Subject: April defended-asset weighting and the 10–12 April window.*

EXHIBITS

Exhibit 1 — The public campaign record, March 2024. Mass strikes on Ukrainian energy infrastructure March 22 and March 29, 2024, with generation (not transmission) increasingly the target class; interceptor

scarcity publicly acknowledged at head-of-state level. Sources: contemporaneous coverage (Reuters, KyivPost, OSW analyses, March 2024).

Exhibit 2 — The sealed input. Public post of September 4, 2023 (item T32/59), preserved at x.com/vijayjyotish/status/1698760242543133031: two three-day windows (March 9–11, April 10–12, 2024), violence-intensity language, no target of any kind named. Advisory documentation, grading, and the record's own stated limits at jyotishintelligence.com/intel/advisory/IA-RU-016. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — Comparative warning leads (as documented in the advisory). Best conventional public warnings ahead of the April 11 strike: a utility CEO's generic grid-vulnerability briefings ~18 days out; national grid-operator blackout guidance ~10 days out; an open-source tempo flag ~2 days out. None calendar-specific. Collection-driven warning is downstream of observable adversary preparation; a calendar-indexed claim is the one input class whose lead is not capped by the adversary's clock — which is also why it arrives unverifiable.

Exhibit 4 — The intake note (composite teaching device). Berg's September note, verbatim in the case text: class logged, limits stated, revisit trigger armed.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · 72-hour reweight	C · Standing reweight
What it adds	nothing	bounded cover where the planner's own model points	permanent posture shift
What it forecloses	the hedge	72 hours of thinner cover elsewhere	flexibility across the whole list
Cost if the window is empty	0	one movement cycle + a dry log line	the saturation adversary's induced error
Cost / value if the strike comes	interceptors standing at standing priorities	a defense at the address the model flagged	right — at unsustainable cost
Speed	—	days, inside tasking authority	doctrine cycle
Authority required	none	allocation planner's own	higher echelon

Reversibility	n/a	by design — the date caps it	slow
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/interceptor-window
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Interceptor Window

Rails: the seal named a violence window, **not a target** — the teaching note must never let "the forecast pointed at Trypilska" survive on the board, because it is false and the case's own protagonist says so; no claim that the strike was preventable; the absolute-scarcity concession is load-bearing and must reach the final synthesis.

LEARNING OBJECTIVES

(1) The reweight as the canonical bounded hedge: how a *dated* window converts "defend everything always" into "defend this, then return," capping the hedge's cost at the window's length; (2) the seam between what a warning says and what the reader infers — and why owning that seam explicitly is the difference between analysis and contamination; (3) scarcity versus allocation: when redistribution is worth its price even though it cures nothing; (4) why calendar-indexed warnings structurally outrun collection-driven ones (Exhibit 3), and what an intake process must therefore do differently.

ASSIGNMENT QUESTIONS

1. Write Berg's Monday memo. Commit to a package and a rationale that survives both worlds.
2. Separate, in writing, what the September seal asserts from what Berg infers. Which half of the April 10–12 recommendation is the hard half, and who owns it?
3. Package B costs 72 hours of thinner coverage elsewhere. Something may be hit *there* instead. Price that world honestly — does Package B still clear?
4. "A reweight redistributes scarcity; it does not manufacture interceptors." Is that an argument against Package B, or merely against overclaiming for it? Where exactly is the line?
5. Design the intake lane: what should a dated, falsifiable, unverifiable-at-receipt input trigger at T-220 days, at T-30, and at T-12?
6. **The machine staff-officer.** An AI fusion cell now proposes daily interceptor allocations, weighting every input it ingests. Feed it dated-window inputs like T32 and it will learn to expect calendars — including ones that aren't there. How should a dated, low-credibility window enter the model (a bounded prior bump? a human-held overlay?), and who owns the decision that the 72 hours have started? Write the targeting-cell standing order.

THE BOARD PLAN (75–90 MIN)

Open on the dispatch analogy and the defended-asset list (10) — engineers and MBAs both own scarcity intuitively; get it on the board before any forecast appears. Build the three packages and price them (20). Introduce Exhibit 2 *with the intake note beside it* (20): run Q2 hard — students reliably inflate the seal into a targeting product, and the case is designed to catch them doing it. Then Exhibit 3 and the structural argument (10). Memo commitment (20), pairs defending A vs. B. Close (10) on Q4: the difference between an input that solves your problem and an input that makes one slice of it — the calendar slice — cheap.

WHAT HAPPENED, AND THE GRADE

On **April 11, 2024** — the middle day of the named window, 220 days after the seal — Russia executed the largest combined-arms strike of the war to that date: 82 munitions in a single overnight wave; the Trypilska Thermal Power Plant, the largest single generator supplying the Kyiv region, was completely destroyed, removing ~1,800 MW of dispatchable capacity in one night. President Zelensky, the same day: *"Russia destroyed Trypilska. Why? Because we ran out of missiles to defend it."* An allocation failure by the defender's own public account — though the record's file, and this note, hold the concession at full weight: scarcity was absolute, and no one can rerun that night with a battery in a different field. The record grades T32 a HIT on its stated terms — a dated violence window, landed on its middle day — and prices its own limits in public: the target inference was never the seal's to sell. Whether Package B would have changed the outcome is unknowable; the ceiling of the claim is recovery-cycle timing, and the case should end at that ceiling, not above it.

EPILOGUE

The strikes that followed — the May–June waves, the 200-munition August wave — turned the April event into the hinge of a campaign, and the question Berg's memo answered into a season-long standing problem. The transfer lesson is the one the cell's junior analyst had bought for the price of a log entry: an intake process that scores inputs by pedigree files the calendar-indexed ones under "unevaluated," and gets 18 days of warning where 220 were on the desk. The reweight posture exists for exactly the class of input that cannot be believed and cannot be ignored — because it is dated, it can be *hedged*, and a dated hedge has a price you can say out loud in a memo that survives the lookback.

CASE SEVEN

The Storm on the Frontage: A Signature in September

A teaching case in operations, heavy-weather readiness, and the rehearsal posture — authority designed before it is needed

> **AUTHORSHIP DISCLOSURE.** This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record. The storm, the strikes, the dates, and the public statements in this case are real and cited. The **protagonist and her terminal are composite characters** — "Daria Ionescu" and the "Pontica Terminal" are not real; their frontage, budgets, and meetings are teaching devices, and no outside-world fact they discuss is invented. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works, and — a rail that must survive every edit — **no position on the storm's causation**: the event that arrived was a meteorological storm, and the case treats it as a storm.
>
> **METHOD DISCLOSURE.** The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

THURSDAY, SEPTEMBER 7, 2023 — 16:00, A WESTERN BLACK SEA PORT

Daria Ionescu ran operations for the Pontica Terminal — a mid-size commercial frontage on the western Black Sea: two berths, a container yard, a grain gallery, a floating crane, pierside stores, and a heavy-weather bill that had last been walked, she was fairly sure, before her promotion. Her Q4 readiness plan was due to the port authority in a week, and Q4 readiness plans were where good intentions went to be deferred: every line item competed against the same winter — dredging, fender renewals, the crane survey — and destructive-weather rehearsal had lost that competition three years running, because the western Black Sea's storms had been survivable and the bill had never actually been executed in anger.

The item that complicated this year's plan had arrived from an unexpected direction: the terminal's insurance broker, in the annual renewal conversation, had forwarded "for whatever it is worth" a public, timestamped

post series of **September 4** — three days old — from a source with a sealed, publicly graded forecast record. The thread stated that a natural-disaster-class sea event — in the words preserved at the source, *"a wave that targets the Crimean coast and the entire Black Sea fleet with it"* — was a live scenario in the forward window, and closed with a posture line: *"Russia must exercise great effort in monitoring the Black Sea."*[1] The source's method, disclosed in its own materials, was Vedic astrology. Its record was public, dated, and graded with the misses kept.

Ionescu's first reaction was the correct one, and she made herself write it down: *this names a coast that is not mine and a fleet that is not mine.* Her terminal sat hundreds of kilometers from Sevastopol. If the input meant anything at all for her frontage, it meant it only at the level of class: same sea, same season, same storm physics — a reminder that the western basin's forgiving run of winters was a run, not a rule.

Her second reaction was the professional one: *no meteorological product on earth can speak to this.* Deterministic forecast skill ends near ten days. Seasonal outlooks return base rates. If a century-class storm were coming to the Black Sea in November, the first instrument capable of seeing it would say so five to ten days out — and five to ten days, at Pontica, was not enough. She knew exactly what her frontage looked like at five days' notice, because she had watched it: the floating crane's move to shelter needed a tug that was usually working, a berth window that was usually sold, and a sign-off from a harbourmaster who was usually asking why this couldn't wait. The bill existed on paper. The authority to execute it did not exist until someone senior said yes, and someone senior had never been asked in calm weather.

THE THREE POSTURES

She sketched what the strange input could actually buy, at its own level — a class warning, nothing more — and priced it for the readiness plan.

Package A — Monitor. File the input; change nothing; rely on the 5–10-day forecast when and if it came. Cost now: zero. Cost in the bad world: executing an unrehearsed bill, on short notice, with authorities still unsigned — the exact configuration in which movable assets don't move.

Package B — Quiet readiness. Three lines in the Q4 plan. *One:* a frontage triage — sort every asset into movable-in-72-hours, movable-with-notice, and fixed; staff days, no permission needed above her own desk. *Two:* walk the heavy-weather bill once before November — find the seized bollard and the tug in

refit while finding them is cheap; a week to schedule, a day to run. *Three:* the crux — a **pre-authorised move order**: asset list attached, tug tasking drafted, harbourmaster's signature obtained now, execution trigger married to the routine short-fuse forecast cue the terminal consumed anyway. The long-lead input would never replace the weather forecast; it would buy the rehearsal and the signature that made the weather forecast *actionable* instead of merely alarming.

Package C — Elevated posture. Move or shelter the movable assets now, ahead of any forecast, on the strength of the seal alone. Cost: real money, sold berth windows, and her credibility — an overreaction to an input of this class, and she struck it through as she wrote it.

Package B's total price was staff hours, one drill, and one awkward conversation with the harbourmaster that would begin, honestly: *"I cannot verify this input, and I am not asking you to believe it. I am asking you to sign, in calm weather, the order you would want already signed in bad weather — because the case for that signature is true whether or not the input is."* That framing mattered to her, because it was the only version of the ask she could defend in either world. If no storm came, the terminal had a rehearsed bill and a standing authority it should have had anyway. If one came, the five usable days of forecast would land on a frontage where the paperwork was already done.

What stood against even that modest ask was the ledger of three deferred years, the renewal season's workload, and the certainty of the question she would face at the port authority review: *"And what is the basis for re-prioritising rehearsal this year?"* — to which the true answer began with an astrologer.

The Q4 plan was due in a week. Ionescu had the triage template open in one window and the readiness plan in the other.

EXHIBITS

Exhibit 1 — The forecasting constraint. Deterministic numerical weather prediction skill ends near ten days; seasonal outlooks return only climatological base rates. On September 7, 2023, no meteorological product could speak to late November. Late-autumn Black Sea storms are a known seasonal feature — the base rate is real, and it is the strongest argument both *for* rehearsal generally and *against* crediting any particular warning.

Exhibit 2 — The sealed input. Public posts of September 4, 2023 (T48–T50/59), preserved at x.com/vijayjyotish/status/1698760261258059811 and adjacent: a natural-disaster-class sea event against the Crimean coast and the Black Sea Fleet, framed in the thread as a "Pearl Harbor type incident in Sevastopol." Advisory documentation, grading, and the record's own stated limits at jyotishintelligence.com/intel/advisory/IA-RU-006 (sealed at 0.78; the record prices the call at ~1-in-375 before the fact — ~8.55 bits — with the partition published for repricing). Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The frontage inventory (composite teaching device). Movable in 72 hours: floating crane (tug + berth window + harbourmaster sign-off), harbour craft, pierside container stores. Movable with notice: grain-gallery spouting, yard reefers. Fixed: berths, breakwater, gallery structure. Prior heavy-weather bill: exists on paper; last walked 3+ years ago; no standing execution authority.

Exhibit 4 — The independent datum. On September 12, 2023 — eight days after the seal — Elon Musk stated publicly that Ukraine had sought Starlink support for a strike on the Black Sea Fleet at anchor in Sevastopol, describing it as "kind of like a Pearl Harbor like attack." Confirmation of the threat *environment* the thread described, from a party with no stake in the seal; not confirmation of any storm.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · Quiet readiness	C · Move assets now
What it adds	nothing	triage + walked bill + signed order awaiting a trigger	motion on the seal alone
What it forecloses	the five usable days	nothing	berth revenue; credibility
Cost if no storm comes	0	staff hours, a drill, one calm conversation	real money, visibly wasted
Cost / value if the forecast window opens	unrehearsed scramble at T-5	execution in the 48–72 hours that exist	already sheltered — by luck
Speed	—	weeks, off-season	days, expensive
Authority required	none	hers + harbourmaster's signature	port authority
Reversibility	n/a	full	n/a

Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/storm-frontage
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Storm on the Frontage

Rails: the match is **scenario-class, not mechanism-class** — the sealed thread framed a deliberate act presenting as a natural disaster, what arrived was a storm, and the case adopts no causal reading whatsoever; the seasonal base rate gets full weight; the seal named a coast and fleet that are *not the protagonist's*, and the generalization from named target to storm class is Ionescu's inference, owned as such.

LEARNING OBJECTIVES

(1) The rehearsal posture as authority design: the insight that the binding constraint at five days' notice is rarely information — it is unsigned paperwork, unwalked procedure, and unpriced tug time; (2) trigger discipline: long-lead scenario warnings convert into *pre-authorisations and rehearsals*, never into immediate movement — their value realises on the day the short-fuse cue arrives; (3) the either-way argument: how to ask an authority for something whose case is true whether or not the input is; (4) reading an input at its own level — class, not address — and owning the generalization explicitly.

ASSIGNMENT QUESTIONS

1. Write the three lines of Ionescu's Q4 plan and the paragraph she uses to defend them at the port authority review — including an honest answer to "what is the basis?"
2. The seal names Crimea and a fleet. Ionescu's terminal is neither. Is her generalization to "same sea, same season, same class" legitimate analysis or contamination? Where is the line, stated as a rule?
3. Package B is nearly free and defensible in both worlds. So why had it lost the budget competition three years running? Name the organizational forces, not the individuals.
4. Draft the pre-authorised move order's trigger clause so a duty officer can execute it alone at 02:00 against a forecast product. What must it name?
5. The harbourmaster signs. No storm comes for five years. Was the signature wasted? Price the answer.
6. **The machine staff-officer.** The terminal's new AI operations system watches every forecast feed and can issue the pre-authorized move order itself the moment trigger conditions are met. What does pre-authorization mean when the executor is a machine — does the harbourmaster's calm-weather signature now authorize an

algorithm, and what trigger-quality standard must a machine-executable clause meet that a duty-officer clause needn't? Rewrite template clause and sign-off accordingly.

THE BOARD PLAN (75 MIN)

Open on Exhibit 3 and the five-day problem (15): have the room walk what actually happens at T-5 with no pre-authorisation — this builds the case for Package B before any forecast, or forecaster, appears. The three packages (15). Introduce Exhibit 2 (15): run Q2 — the room should discover that the input's only legitimate use for *this* desk is as a class reminder, and that the strongest version of Package B barely depends on the input at all. That is not a weakness; it is objective 3, and the instructor should name it. The review-board role-play (20): one student defends the plan, one chairs, the room asks "the basis" question. Close (10) on Q5 — readiness as an asset with option value, not a cost with no event.

WHAT HAPPENED, AND THE GRADE

On **November 26–27, 2023** — 83 days after the seal, some 73 of them beyond the reach of any forecast product — the most powerful Black Sea storm in decades struck the basin: waves near nine meters, roughly 1.9 million people without power across the region, coastal infrastructure washed out. At the named coast, Black Sea Fleet shore infrastructure at Sevastopol was wrecked, floating assets moored on the exposed frontage — including the fleet's known dolphin pens — were destroyed, and, compounded by a parallel kinetic campaign (the fleet HQ struck 18 days after the seal; a landing ship destroyed at Feodosia in December), the fleet's principal Crimean basing became untenable and major assets redeployed east through 2024. The record grades the September call a HIT **on the sealed scenario's falsifiable elements only** — a destructive wave event against the named coast and fleet, in window — and volunteers the two concessions this note enforces: the mechanism did not match the thread's framing, and autumn Black Sea storms sit on a real base rate; the priced improbability lives in the conjunction (century-class, named coast, named fleet, sealed window), and the record hands readers the partition to reprice it. For Ionescu's composite frontage the storm was survivable weather — and the teaching point is that she could not have known which side of the basin the bad conjunction would land on, which is exactly why the posture she was pricing is called *rehearsal* and not *prediction*.

EPILOGUE

The question that outlives the case is the one the paper behind it leaves on every planner's desk for Monday morning: *of everything on my frontage, what can move in 72 hours — and is the order to move it already signed?* A terminal that can answer both halves has spent staff hours and one calm-weather conversation. A terminal that cannot has decided, by deferral, to hold an unpriced option short — and the Black Sea, on the record of November 2023, sometimes exercises.

CASE EIGHT

The August Package: Continuity Before the Cause

A teaching case in public administration, emergency authorities, and drafting the package before the disaster names itself

> **AUTHORSHIP DISCLOSURE.** This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record. Hurricane Helene, its toll, and North Carolina's documented election-administration response are real and cited; North Carolina's response was by any professional standard a success, and this case depicts none of its internal deliberations. The **protagonist and his agency are composite characters** — "Marcus Hale" and his unnamed Southeastern coastal state are not real; their board, statutes, and meetings are teaching devices, and no outside-world fact they discuss is invented. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works.
>
> **METHOD DISCLOSURE.** The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

TUESDAY, SEPTEMBER 3, 2024 — 08:45, A SOUTHEASTERN STATE CAPITAL

Marcus Hale had been executive director of his state's board of elections long enough to know the shape of the next sixty days by heart. The lull he was standing in — the last quiet week before early-voting preparation consumed every senior staff hour — was the only slack in the calendar between now and certification. Whatever was not drafted, cleared, and signed by the end of September would, if needed later, be drafted in a crisis instead.

The item he could not decide whether to take seriously had come through the informal channel that election directors actually use: a colleague's note on the interstate listserv, half-joking in tone, flagging a publicly timestamped post of **August 27** about the neighboring swing state. Alongside an electoral call that was none of Hale's business, the post carried one operational sentence, preserved verbatim at the source: *"There is a potential for severe weather in early November. Even though it is the end of hurricane season... Such weather could cause*

delays and interruptions in the vote counting process."[1] The source's method, disclosed in its own materials, was Vedic astrology; its record was public, sealed, and graded with misses retained. The colleague's gloss: *"Your hurricane exposure is worse than ours. Do you have a package?"*

Hale knew the honest answer, which was why the joke had not landed as one. His board, like most, held broad statutory authority for election-continuity measures — site relocation, absentee accommodations, plan modifications — exercisable by board vote. What it did not hold was anything *pre-drafted*: no activation trigger, no template resolution citing the existing authority, no counsel-cleared menu of relief categories. If a disaster hit his state's election infrastructure, the package would be built the way packages are always built after landfall — in about ten days, by exhausted senior staff, at the exact peak of early-voting preparation, under press scrutiny, with counsel reviewing at crisis speed the same clauses that could have been reviewed at leisure in September.

THE CLAUSE, READ AT ITS OWN LEVEL

Hale forced the discipline his job required. The sentence named a weather-event class and one effect — delays and interruptions in the vote *counting* process. It named no storm; the National Hurricane Center had nothing on its charts.[2] It named no county, no mechanism, no remedy. And it named a window — early November — that was, he noted, the *least* likely part on climatology: late-season landfalls in his region were rare, which cut both ways, since a severe-weather caveat for a Southeastern coastal state during hurricane season also sat on a real seasonal base rate. A skeptic on his own staff would call it an inexpensive hedge, and the skeptic would have a point.

But the decision the clause pointed at, he kept noticing, did not actually depend on the clause being right. A generic, **cause-agnostic** continuity package — trigger, template resolution, counsel-cleared relief menu — served any weather-driven disruption in any month, and served it identically whether the prompt had been an astrologer, an insurance renewal, or a staff retreat. The input's only real function was to put the question on his desk during the one week when answering it was cheap.

Package A — Monitor. File the note. Rely on the board's raw authority if the season produced something. Cost now: zero. Cost in the bad world: the ten days.

Package B — The pre-drafted menu. Three staff-days and one counsel review, this month: activation trigger tied to a declared emergency touching election infrastructure; template emergency resolution citing chapter and verse of existing authority; a menu of relief categories (absentee-channel accommodations, site relocation, cross-jurisdiction options, plan modifications) each pre-annotated with its legal basis and its litigation exposure. Publishable in advance — which mattered, because improvised late changes to voting procedure were a documented surface for disinformation, and a package announced in September is a package no one can call a November power grab.[3]

Package C — The elevated posture. Stand up a weather-contingency working group, pre-position county-level plans, run an exercise. Real staff cost at the worst time; defensible only if he was prepared to treat one unverifiable sentence as tasking, which he was not.

The deputy's objection at the morning meeting was the institutional one: *"Counsel's queue is full, the counties are screaming about poll-worker recruitment, and you want three senior staff-days on a hypothetical because of a tweet about the neighbors?"* The counter, which Hale tried out loud: *"The package is the cheapest item on our whole preparedness menu, it is cause-agnostic, and the only thing the tweet changes is* when *we draft it — now, at leisure, or later, in the ten worst days of our year."*

Counsel could take the review this week or not until October. Hale had to decide by Friday.

EXHIBITS

Exhibit 1 — The operating environment, August 27, 2024. NHC charts clear; no named system threatening the Southeast. Polling in the neighboring state effectively tied (48–48 in the in-state academic poll). No consensus forecasting product — polling averages, academic models, in-state polls — carried severe weather as an election-administration axis.

Exhibit 2 — The sealed input. Public post of August 27, 2024, 11:55 ET (x.com/vijayjyotishusa/status/1828506539860336747): the severe-weather clause quoted verbatim in the case text, embedded in a broader sealed forecast for North Carolina. Advisory documentation and grading — including the grades that went against it — at

jyotishintelligence.com/intel/advisory/IA-US24-008. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — What the ten days look like (documented, neighboring state). After Helene devastated western North Carolina (Sept 27–28: ~107 deaths per NCDHHS, 6,000+ displaced, mail service suspended at dozens of post offices), the NC State Board of Elections signaled emergency planning Sept 30 and adopted a unanimous, bipartisan emergency-measures package Oct 7 for what became 25 disaster counties — drafted, cleared, and voted in roughly ten days, at the peak of early-voting preparation. Early voting opened on schedule with 76 of 80 disaster-county sites on day one and a record first day statewide. Sources: NCSBE press releases of Sept 30 and Oct 7, 2024; TIME, Oct 2, 2024; Brennan Center, Oct 20, 2024.

Exhibit 4 — The package skeleton (composite teaching device). One page: trigger clause; template resolution; relief-category menu with statutory citations and litigation notes; publication plan.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · Pre-drafted continuity menu	C · Cross-state elevated posture
What it adds	nothing	trigger + template resolution + counsel-cleared menu	working groups, exercises
What it forecloses	the September lull	nothing	staff at the worst time
Cost if the season is quiet	0	three staff-days, published and waiting	real drag on prep
Cost / value if disaster hits the calendar	ten days of crisis drafting at the peak	activation instead of drafting	partially ready, exhausted
Speed	—	days, this month	months
Authority required	none	his own + counsel hours	board + partners
Reversibility	n/a	full	n/a
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/august-package
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The August Package

Rails: North Carolina's real response was a *success* and is never graded here — the counterfactual is margin and tempo, not rescue; the sealed clause named early November while Helene struck late September — a three-week timing miss the record itself grades PARTIAL, and a companion sub-claim in the same post graded MISS; both concessions are load-bearing. The seasonal base rate gets full weight.

LEARNING OBJECTIVES

(1) The cause-agnostic package as the purest form of the rehearsal posture: readiness that does not require the warning to be right, only the question to arrive early; (2) tempo as the priced asset in public administration — the difference between crisis *drafting* and crisis *activation*, paid in senior staff hours at their scarcest; (3) the legitimacy dividend: pre-published authorities as a defense against the disinformation surface that improvised changes create; (4) how to justify low-cost preparedness triggered by a low-credibility input — the "the case for the signature is true either way" argument, and its limits.

ASSIGNMENT QUESTIONS

1. Write Hale's Friday decision memo. If Package B, write the paragraph that answers "why now?" without hiding the input's source; if Package A, defend the ten days.
2. The clause missed its own window by three weeks (named early November; the storm came in late September). A director who staged to the *calendar* rather than the *axis* might have been caught staged for the wrong month. What does this teach about which part of a warning to consume?
3. The record grades the weather clause PARTIAL and a companion sub-claim in the same post MISS. Does an input's mixed grade change what Package B was worth? Should it have changed Hale's Friday decision, made before any grade existed?
4. Exhibit 3 shows the neighboring state succeeding *without* a pre-drafted package. Steelman the conclusion "so the package wasn't needed" — then find what that argument prices at zero.
5. Design the standing rule: what class of unverifiable operational flag should reach a state election director's desk, and what is the maximum action it may trigger?
6. **The machine staff-officer.** An AI legal-drafting tool can generate the entire continuity package — trigger, resolution, relief menu with citations — overnight, in every state, every year. Which of the ten days does machine drafting actually compress, and which clocks (counsel judgment, board vote, pre-publication

legitimacy) does it not touch? Then the harder question: does effortless drafting erode the pre-commitment value that made the package trustworthy? Write the adoption memo.

THE BOARD PLAN (75 MIN)

Open on the calendar (10): the September lull versus the October peak — students must feel that the *same* three staff-days have two radically different prices six weeks apart, because tempo is the case's hidden currency. The three packages (15). Introduce Exhibit 2 with its grades' eventual existence concealed (10) and force Friday's decision (20) — memo written, positions defended. Then reveal Exhibit 3 and the grades (10): run Q2 and Q4 — the room should discover that Package B's value survived the timing miss *because Hale consumed the axis, not the date*, and survived the neighbor's success because ten days of senior-staff crisis drafting is a real cost even when the outcome holds. Close (10) on Q5 and the legitimacy dividend.

WHAT HAPPENED, AND THE GRADE

Tropical Depression Nine formed on September 24, 2024 — 28 days after the sealed post. As Hurricane Helene, it devastated western North Carolina 40 days before the election: roughly 107 deaths in-state, thousands displaced, postal service suspended at dozens of post offices — a direct hit on absentee infrastructure. The neighboring state's board drafted and unanimously adopted its emergency package in roughly ten days and ran a successful disaster-shadowed election, with disaster-area turnout outpacing the rest of the state. The record grades the sealed clause **PARTIAL** — the axis (severe weather disrupting vote administration, in-state) materialized with unusual force; the named window was wrong by three weeks — and grades the same post's Cape Fear sub-claim **MISS**, both retained at full weight. The teaching synthesis: the input earned consideration as a *class*, not deference as an oracle; the posture it pointed at was buildable from its own words, cost three staff-days, and was worth owning in every world the season could produce — including the one where nothing happened and the package waited, published and legitimate, for a future year's storm.

EPILOGUE

The doctrine North Carolina's executive director stated publicly during the real crisis — *"We do not stop an election; we figure out how to proceed"* — is the sentence every continuity package exists to make cheap. The case's residue for a public administrator is a single off-season question, cousin to the port planner's: *of everything my agency would need signed on the tenth day after a disaster, how much could counsel clear this month — and what, exactly, is stopping me from asking?*

CASE NINE

The Swap-Equity Map: Building the File Before the Phone Rings

A teaching case in corporate security, duty of care, and the package posture — assembling the deliverable before the crisis requests it

> **AUTHORSHIP DISCLOSURE.** This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record. The detentions, trades, dates, and public statements in this case are real and cited; the case depicts no real news organization's internal deliberations, makes **no claim that any detention could have been prevented**, and treats every detained journalist named in it with the seriousness their ordeals demand. The **protagonist and her organization are composite characters** — "Claire Duval" and "Interpress" are not real; their bureaus, budgets, and insurance schedules are teaching devices, and no outside-world fact they discuss is invented. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works.
>
> **METHOD DISCLOSURE.** The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

TUESDAY, MAY 2, 2023 — 18:20, LONDON

Claire Duval ran global security for Interpress, an international wire with correspondents in forty countries and, at this moment, three of them in Moscow. For thirty-three days — since FSB officers arrested the Wall Street Journal's Evan Gershkovich in Yekaterinburg on espionage charges — her job had contained a question that the entire industry was asking and no one could answer: *what is the actual threat model for a Western journalist in Russia now, and what does a responsible employer change?*

The industry's public answer was advocacy. The Committee to Protect Journalists and Reporters Without Borders were running the frame the moment demanded — condemn the detention, demand immediate unconditional release.[1] Duval endorsed every word of it and knew it was not a security posture. Advocacy is what you say about the detained. Her job was

the undetained: the three passports in Moscow, the insurance schedule renewing in June, and the standing question of what Interpress would do in the first seventy-two hours if the next name in a Russian courtroom was one of hers.

The odd item in her intake had been forwarded by her deputy with the subject line *"price this?"* A public, timestamped post of **April 30**, from a source with a sealed and publicly graded forecast record, stated — in the words preserved at the source — that Russia would *"be able to successfully negotiate a mutual exchange of POWs and others held for espionage/illegal activities,"* and that *"the WSJ reporter Evan Gershkovic might see relief during this period."*[2] A named journalist, and a named mechanism: a prisoner swap. Not a diplomatic release, not a pardon, not a collapse of the charges — an exchange. The source's method, disclosed in its own materials, was Vedic astrology; its ledger kept its misses at full weight.

Duval read it the way she read everything: for what it asserted, and for what acting on it would cost. The Gershkovich clause was hedged — *might* see relief — and carried no window. The firm clause was the mechanism. And the mechanism, if she priced it as live, reframed her entire portfolio: it meant Russia's operating logic toward Western detainees was hostage diplomacy — detainees as inventory, accumulated for exchange — in which case the protective variable in her Moscow staffing was not the press card. It was the passport.

THE FOUR MOVES, PRICED

That evening she drafted what a posture would actually contain, because the security review was Thursday and "monitor the situation" was not a deliverable.

Passport-based staffing. Rotate US-passport correspondents to non-Russia assignments; weight in-country datelines toward staff whose passports carried less exchange value. Cost: a real coverage compromise, some rotation expense, and a precedent the editors would fight. Speed: weeks. Authority: hers and the bureau chief's.

The rider, bound before the arrest. Interpress's kidnap-and-ransom cover, like most, was drafted around non-state actors and ransom. Whether it responded at all to a *sovereign* holding an employee for exchange was a question she realized she could not answer from memory — and the structural feature of all insurance made the timing brutal: a detained

correspondent is an uninsurable one. The review-and-bind decision existed only until the next arrest. Cost: a premium she'd have to get quoted. Speed: inside the June renewal. Authority: hers, with the CFO's signature.

The swap-equity map. The move that felt least like corporate security and most like the thing the moment actually required. If the mechanism was exchange, the eventual deal would be assembled from a specific commodity: Russian operatives in Western custody whom Moscow wanted back. The asking price was already legible in public — when Moscow excluded Paul Whelan from the December 2022 Griner–Bout trade, reporting identified the name it wanted instead: Vadim Krasikov, the FSB officer serving a life sentence in Germany for the Tiergarten assassination.[3] A standing dossier — which Russian nationals sit in US, German, and allied custody; which have intelligence value to Moscow; which governments hold them and what each has historically required to trade — was analyst-weeks of open-source work, entirely inside her budget. Interpress could never conduct that diplomacy. But if one of hers was ever taken, the difference between arriving at the US hostage-affairs envoy's office on day two with the file built, versus month two with a blank page, was the difference her function existed to buy.

Pre-staged designation package. The legal trigger for the US hostage-recovery apparatus was the Secretary of State's "wrongfully detained" determination — twelve days, in Gershkovich's case.[4] A pre-assembled package per correspondent — employment records, assignment history, accreditation chain — plus standing contacts at the envoy's office, might compress that to days. Cost: phone calls and a template. Authority: nobody's permission required at all.

Against the bundle stood the objection her deputy had already made, because it was fair: *by May 2023, wasn't "swap, eventually" becoming the obvious outcome anyway?* The Reed and Griner trades had built the channel; State had designated Gershkovich within twenty days; within two weeks CNN would report Washington weighing Russian spies as trade material.[5] If the room was drifting toward the same conclusion, what exactly did the strange post add — and if the answer was "nothing," did that subtract from the four moves, every one of which was justified by the *mechanism being live* regardless of who had named it first?

Thursday's review needed a recommendation: the bundle, parts of it, or the industry's actual posture, which was — on the public evidence — none of it.

EXHIBITS

Exhibit 1 — The public record, spring 2023. March 29: Gershkovich arrested in Yekaterinburg. April 10: State Department wrongful-detention designation (12 days post-arrest). Industry frame: unconditional-release advocacy (CPJ, RSF); no outlet or Russia-program shop publicly specified a prisoner-swap mechanism for Gershkovich as of April 30. Prior trades: Trevor Reed (April 2022, for a convicted trafficker); Brittney Griner (December 2022, for arms dealer Viktor Bout — Whelan excluded, Krasikov signalled as the price).

Exhibit 2 — The sealed input. Public post of April 30, 2023 (T19/25), preserved at x.com/vijayjyotish/status/1652910740787871750: the mutual-exchange mechanism stated firmly; the Gershkovich relief clause hedged ("might"), unbounded. Advisory documentation and grading at jyotishintelligence.com/intel/advisory/IA-RU-017. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The insurance gap (composite teaching device). Interpress K&R policy: responds to kidnap by non-state actors, ransom, extortion; state-detention response ambiguous; rider quotable at renewal; binding condition — no known circumstance at inception. Renewal date: June 30, 2023.

Exhibit 4 — The asking-price ladder (public record). April 2022: Reed ↔ convicted trafficker. December 2022: Griner ↔ Bout (arms dealer); Whelan excluded. Signalled next price: a serving FSB assassin in allied custody. Each completed trade rebased the next demand.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · Quiet readiness bundle	C · Elevated posture
What it adds	nothing	staffing map + bound rider + equity dossier + pre-staged designation package	public posture shifts
What it forecloses	the bind-before-event windows	coverage compromises, rotation costs	newsroom independence optics
Cost if no one is taken	0	premium + analyst-weeks + awkward editor meetings	credibility with correspondents
Cost / value if the next arrest comes	day two starts at a blank page	day two starts at the file	visibly prepared — and watched

Speed	—	weeks (map), renewal cycle (rider)	months
Authority required	none	hers + CFO for the premium	editor-in-chief
Reversibility	n/a	full	poor
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/swap-equity-map
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Swap-Equity Map

Rails: no prevention claims — Alsu Kurmasheva's October 2023 detention occurred during personal travel and no employer posture is claimed to have prevented it; the hedged, unbounded character of the Gershkovich clause is load-bearing; the "modal outcome anyway" objection gets its full force and the note's synthesis must run *through* it, not around it.

LEARNING OBJECTIVES

(1) The package posture: when a warning names a *mechanism*, the highest-value response is often to pre-assemble the deliverable that mechanism will eventually demand — the leverage file, the designation package — so the crisis starts at execution, not research; (2) duty of care under a re-based threat model: what an employer owes the undetained when detention risk becomes sovereign and instrumental; (3) bind-before-the-event structures (insurance, designations) and how warning value concentrates on decisions with closing windows; (4) the "obvious in retrospect" trap: distinguishing what the room was drifting toward from what any desk had actually operationalized.

ASSIGNMENT QUESTIONS

1. Write Duval's Thursday recommendation. Which of the four moves clears, at what trigger, and how is it framed to editors who will read passport-based staffing as capitulation?
2. The deputy's objection: "swap-eventually was becoming the modal outcome anyway." Grant it fully. Which of the four moves lose their justification? (Expected answer: none — then ask what that reveals about where the input's value actually lived.)
3. The swap-equity map costs analyst-weeks and will very likely never be used. Price it as an option. What is the strike event, and what does the premium buy on day two?
4. The rider decision exists only until the next arrest. Construct the class of decisions with this structure across other industries — and the intake rule that flags them.
5. A named mechanism, hedged subject, no window. Draft the one-paragraph intake note that neither over-reads nor under-reads this input.
6. **The machine staff-officer.** An AI OSINT system now maintains the swap-equity map continuously — every Russian national in allied custody, every signaled asking price, refreshed daily. The analyst-weeks are gone. What did building the

map by hand teach the desk that consuming a dashboard doesn't, which uses of the map (SPEHA day two, the rider bind decision) remain irreducibly human, and what's the audit that catches the day the machine's map quietly goes stale? Write the standing order.

THE BOARD PLAN (75–90 MIN)

Open on the duty-of-care frame (10): three passports in Moscow, and what "monitor" means to the people carrying them. Build the four moves and price each (25) — insist on cost, speed, and *authority required*; students discover all four sit inside the security function's own discretion, which is the quiet architecture of the case. Introduce Exhibit 2 (15): run Q5, then Q2 hard — the room should split between "the post added nothing" and "the post was the only dated, named-mechanism statement anyone had," and the instructor's move is to show both camps converging on the same Thursday recommendation, which is objective 4 made visible. The insurance window (10): Q4, generalized. Commit (15): the recommendation memo. Close (10) on the package posture's general form.

WHAT HAPPENED, AND THE GRADE

The mechanism held. On October 18, 2023, Russia detained a second Western journalist — RFE/RL's Alsu Kurmasheva, in Kazan — doubling the journalist-class inventory. On February 2, 2024, Kremlin spokesman Dmitry Peskov publicly confirmed the silent swap channel. On July 19, 2024, a closed court sentenced Gershkovich to sixteen years — the conviction that, in Russian practice, formalizes a detainee as tradable. Thirteen days later, on **August 1, 2024** — 460 days after the seal — Gershkovich, Whelan, and Kurmasheva walked free at Ankara in the largest East–West exchange since the Cold War: twenty-four prisoners, seven countries, and at its center exactly the asking price the ladder had signalled — Krasikov, released by Germany. The record grades the April 30 statement a **HIT** on the mechanism, and its own file concedes what this note enforces: the subject clause was hedged and unbounded, a 460-day materialization is a long leash, and "swap, eventually" had a rising base rate after Reed and Griner. What the base-rate account does not explain is the operational silence: the organizations paid to make the inference did not publicly make it, on that date, for that man, by that mechanism — and on the public evidence, no desk ran the four moves until the process forced them. The value of the input, as always in this book, was not that it was right. It was that it arrived dated, named, and priceable, on a day when everything it implied was still cheap.

EPILOGUE

Held: 491 days (Gershkovich), 288 (Kurmasheva), 2,043 (Whelan) — the denominators any tempo gain would have worked against, and the numbers that should silence any classroom impulse to treat this case as a puzzle with a clever answer. The package

posture's residue for a security director is one standing question: *for each mechanism that could take one of my people, does the file that mechanism will eventually demand already exist — and whose day two does it save?*

CASE TEN

The Trade-Profile Question: Two Lines on a Proposal Form

A teaching case in underwriting strategy, the reprice posture, and acting on a warning of which half will turn out to be wrong

> **AUTHORSHIP DISCLOSURE.** This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the sealed input in this case is graded **PARTIAL** on that record, with two of its four named vectors marked as failures, a fact the case is built around rather than away from. The vessel losses, dates, and attributions cited are real and sourced. The **protagonist and her syndicate are composite characters** — "Priya Raman" and the "Halcyon Syndicate" are not real; their book, rating model, and meetings are teaching devices, and no outside-world fact they discuss is invented. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works, and nothing in it is underwriting advice.
>
> **METHOD DISCLOSURE.** The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

MONDAY, SEPTEMBER 11, 2023 — 10:30, LONDON, EC3

Priya Raman underwrote hull war and cargo war for the Halcyon Syndicate, and her September problem was the usual one: the January renewals were being shaped now, and her rating model knew exactly one big thing about the Russia–Ukraine war. It knew *where* it was. War risk priced by geography — listed areas, breach premiums, the combat littoral. A vessel entering the northwestern Black Sea paid; the same vessel in the mid-Mediterranean did not, because the mid-Mediterranean was not where wartime losses happened. The model was not wrong. It was the distilled experience of every maritime war in living memory: ships are lost where the fighting is.

The item her cargo analyst had flagged at the Monday meeting — with the preamble "you'll laugh, but read it to the end" — cut against that one big thing. A public post of **September 4**, timestamped on the platform, from a source maintaining a sealed and publicly graded forecast record, stated verbatim: *"Hijack and/or destruction of grain and supply ships are possible, especially in*

international waters. Hijack of airliners, ships and facilities with Russians could be used to increase pressure on the honorable President Putin."[1] The source's method, disclosed in its own materials, was Vedic astrology.

Raman did what a class underwriter does with any submission: she decomposed it. Four distinguishable vectors. *Destruction of supply ships. Grain vessels as a named class. A hijack mode.* And the qualifier that made her stop laughing — *"especially in international waters."* In September 2023 that was an unusual thing for anyone to write down. Russia had exited the grain deal in July; the maritime war was concentrating in the littoral, exactly where her model priced it; high-seas losses were part of no pattern anyone in the market was rating. Three of the four vectors read like a hedged spectrum. The fourth read like a claim about *where her model was blind*.

THE DECOMPOSITION, AND THE SCREEN

She was clear with herself about what was not on the table. Nobody reprices a book on a single public post — not from a bank's research desk, and certainly not from this. If the input had any use at all, it had to point at actions whose cost survived the input being mostly wrong. She listed three.

Two proposal-form questions. *Does the vessel call at Russian ports? Does the trade serve Russian export logistics?* Her forms did not ask, because her model did not care — geography was the rating variable, and a tanker's customers were its own business. Adding the questions cost nothing, sat entirely within her authority, and began building the one dataset the market did not have: her book's exposure by *trade profile* rather than by position.

A referral rule. Russia-serving trades flagged out of the fleet-terms flow for individual review. Not a rate change — a routing change. Cost: some friction with brokers, a few hours a month.

A wording review. If a loss of contested origin ever arrived *outside* the listed areas — an "onboard explosion" that investigators later called a mine — did it attach to war perils or marine perils? That question was far cheaper to think through at renewal, in daylight, than in litigation after the fact. One session with claims counsel.

Total price of the screen: near zero, one renewal cycle, no reinsurance conversation, no capital action. Against it, the objections her deputy underwriter raised at the Thursday review, all fair. *The spectrum problem:* if they acted on the whole post, they'd also be loading grain trades and seizure covers — and if half the spectrum never fired, the syndicate would have

chased phantoms on an astrologer's say-so; how did she propose choosing which half to believe? *The tempo problem:* even if something eventually happened in international waters, "eventually" might sit outside any single treaty year — what was the January book actually buying? *The trend problem:* Ukrainian strikes on Russian maritime assets were an established campaign; if losses ever spread beyond the littoral, wasn't that just the trend extending, foreseeable without anyone's post?

Raman's answer to the first objection was the design of the screen itself: she was not choosing which half to believe. Questions and referral rules pay for themselves even when most of the spectrum dies, because their function is not prediction — it is making the book *legible along a new axis* before the market rates that axis. Her answer to the third was the qualifier: trend-following would have written "more losses near the war." Someone had written *"especially in international waters"* while every observable loss was landing somewhere else.

The January forms went to print in three weeks. The Thursday review needed her decision on all three items — and one sentence, for the file, on why.

EXHIBITS

Exhibit 1 — The market's one big thing, September 2023. War-risk rating by listed areas and breach premiums; wartime losses concentrated in the Black Sea littoral; grain-corridor collapse (July 2023) the market's active repricing event; high-seas sabotage of merchant tonnage part of no priced pattern.

Exhibit 2 — The sealed input. Public post of September 4, 2023 (17/59 in the thread), preserved at x.com/vijayjyotish/status/1698760227435221469, quoted verbatim in the case text. Advisory documentation and grading — **PARTIAL**, under the record's named-spectrum rule: supply-ship destruction and the international-waters axis HIT; grain-specific and hijack vectors MISS — at jyotishintelligence.com/intel/advisory/IA-RU-030. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The screen, priced (composite teaching device). Item 1: two proposal-form questions — cost ≈ 0, authority: class underwriter. Item 2: referral rule for Russia-serving trades — cost: broker friction, hours/month. Item 3: war-vs-marine-perils wording review for contested-attribution losses outside listed areas — cost: one counsel session. Total: inside one renewal cycle; no reinsurance or capital action.

Exhibit 4 — What subsequently arrived (public record, for the teaching note). Jan 21, 2024: Ust-Luga fuel terminal fire after reported drone strike. Feb–Mar 2024: *Ivanovets*, *Tsezar Kunikov*, *Sergey Kotov* sunk by sea drones (littoral). Aug 22, 2024 — 353 days after the seal: rail ferry with thirty fuel tanks struck and sunk at Port Kavkaz. Dec 24, 2024: *Ursa Major* sinks off Spain after onboard explosions. Feb 15, 2025: tanker *Seajewel* limpet-mined off Savona, Italy — part of a campaign of blasts on tankers serving Russian ports, in international and foreign waters, attributed by investigators to state-actor sabotage. Explicitly excluded: the Dec 2024 *Volgoneft* break-ups (storm losses, not attacks).

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · The screen	C · Reprice the book
What it adds	nothing	trade-profile data + referral lane + wording clarity	rate on an unpriced axis
What it forecloses	legibility on the new axis	broker friction, hours	competitiveness if the axis is dead
Cost if the spectrum dies	0	near zero — the screen pays either way	mispriced book, lost accounts
Cost / value if international waters fire	claims land on a blind book	the book is legible before the market reprices	right — early, alone, unprovable
Speed	—	one renewal cycle	reinsurance conversation
Authority required	none	class underwriter's own	committee
Reversibility	n/a	full	slow
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/trade-profile
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Trade-Profile Question

Rails: the PARTIAL grade is the case's spine — two vectors fired, two died, and any telling that flattens this into "the forecast was right" has broken the case; the tempo concession (the high-seas axis fired past a single policy year) and the trend-following objection both get full weight; no underwriting-loss figures exist for the sabotage losses and none may be invented.

LEARNING OBJECTIVES

(1) The reprice posture at its cheapest: before changing a rate, change what you *ask* — a screening question converts an unpriceable warning into a dataset; (2) spectrum discipline: how to act on a multi-vector warning without betting on any vector — choosing actions that pay even when half the spectrum dies; (3) rating-model blindness: recognizing when an input's value is not its prediction but its *axis* (trade profile vs. geography), and what it costs to become legible along a new axis after the market does; (4) grading honesty: what a PARTIAL means, why a record that marks its own calls down is the only kind worth reading, and how a decision made before the grade existed should be evaluated.

ASSIGNMENT QUESTIONS

1. Write Raman's one sentence for the file — the rationale that justifies the screen without endorsing the source, survivable in the world where nothing ever happens in international waters.
2. Decompose the sealed post into its four vectors and assign each a screen-level action or a deliberate non-action. Defend the asymmetry.
3. The deputy: "If losses spread beyond the littoral, that's just the trend extending." What, precisely, does the trend predict that the post doesn't — and vice versa? Which difference is ratable?
4. The high-seas axis fired 15–17 months after the seal — outside the January treaty year. Does that void the screen's value for the January book? What does it say about which *instruments* (annual rating vs. wordings vs. data collection) each part of a long-lead warning should feed?
5. It is February 2025 and the *Seajewel* claim lands on a competitor whose form never asked the trade-profile question. Reconstruct their position — attachment, data, negotiation — versus Halcyon's.

6. **The machine staff-officer.** An AI underwriting copilot can score every submission for Russia-serving trade exposure from raw AIS history — no proposal-form question needed. But an unaudited scoring model is a rate change nobody approved. Where does the copilot sit relative to the referral rule, what documentation preserves the wording-review discipline when a machine flags the risk, and who owns the day the model's 'trade profile' definition drifts? Write the governance paragraph.

THE BOARD PLAN (75 MIN)

Open on the rating model (10): make the class own "geography is the variable" as *correct historical inference*, not stupidity — the case has no villains, only a model meeting a new war. Decompose the post on the board (15): Q2 — the four vectors, and the qualifier as the only genuinely surprising content. Build the screen and price it (15). The Thursday objections (15): run Q3 and Q4 against the room; the spectrum objection should be answered by the screen's design, not by faith in the source. Commit (10): the file sentence, written. Reveal Exhibit 4 and the PARTIAL (10) — run Q5, then close on objective 4: the decision was right or wrong on September 11, at the prices of September 11, and the grade that arrived later changes neither — it only tells you the record was honest enough to keep score.

WHAT HAPPENED, AND THE GRADE

The littoral campaign continued; the logistics chain was hit at Ust-Luga and Port Kavkaz; and from December 2024 the strangest specification fired hardest — tankers serving Russian ports began suffering blasts in international and foreign waters, the *Ursa Major* and the limpet-mined *Seajewel* among them, attributed by investigators to state-actor sabotage. The market's rating axis had moved from geography toward trade profile, exactly the direction the two proposal-form questions had pointed. The record grades the September 4 post **PARTIAL** and states its failures as plainly as its successes: no grain vessel was ever targeted; no hijack of any kind occurred; a desk that had acted on the full spectrum would have loaded covers against losses that never came. That is why the case's recommended posture was a screen and not a rate change — and why the input's surviving value concentrated where the surprise was: in the five words about international waters, written while every loss in the world was landing somewhere else.

EPILOGUE

Two lines on a proposal form are the smallest reprice in this book, and the purest statement of its economics: the cost of asking was zero, the cost of not asking was compounding invisibly with every Russia-serving risk that flowed through at fleet terms, and the difference between the two books only became visible on the day a contested explosion off a foreign coast turned into a claim. The residue for any underwriter — of ships, cyber, credit, anything — is a standing question about the model itself: *which variable does my form not ask about, because the last war didn't need it?*

CASE ELEVEN

Two Readers, One Input: The Window That Meant Nothing and Everything

A teaching case in information economics — why the same warning can be decision-relevant at one desk and unusable at the next, with both desks right

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and this case is built on the record's own published concession that part of its most striking-looking success was a **scheduling coincidence.** The launch, the failure, the figures, and the FAA action are real and cited; Blue Origin and AST SpaceMobile appear only through their public record, no villain is depicted, and space is hard. The **protagonists and their firm are composite characters** — "Elif Kaya," "Dr. Sam Whitaker," and the "Aldgate Space Risk Agency" are not real; their binder participation and meetings are teaching devices, and no outside-world fact they discuss is invented. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

TUESDAY, APRIL 21, 2026 — 09:00, LONDON

Two days after the loss, the managing director of Aldgate Space Risk — a small agency underwriting launch and in-orbit covers — put one question to the two people in the room and asked for a joint memo by Friday: *"Do we, or do we not, build a formal intake lane for this input class on future first-reflight binders?"*

The loss itself was public arithmetic. On April 19, Blue Origin's New Glenn had flown its NG-3 mission — a reused booster, its first paying reflight customer — carrying AST SpaceMobile's Bluebird Block 2 FM2, insured for $30 million. The countdown held at T-3:57; after a forty-minute recycle, the vehicle lifted off at 07:25 EDT. The booster separated and landed cleanly on the droneship. Then the upper stage's second burn produced insufficient thrust; the satellite came off at roughly 154 × 494 km against an intended

circular orbit, too low to sustain operations, and was declared a total loss recoverable under its policy. The FAA grounded the vehicle pending investigation.[1] Aldgate carried a line on the cover. The claim would be paid; that was what the product was for.

The reason for the Friday memo was the document that had been sitting, unpriced, in the public domain since the afternoon before the flight.

THE INPUT

On **April 18**, a sealed public advisory — from a source whose method, disclosed in its own first lines, was Vedic astrology, and whose record was published, graded, and Bitcoin-anchored with misses retained — had partitioned the entire April 19 launch day into ten named windows and assessed the day as a whole as **no-go.** Window 2, a nine-minute slot from 07:18 to 07:27 EDT, was labeled *"high risk of catastrophic outcome."* Five anomaly categories were named in advance, the fifth reading in part: *"the Bluebird Block 2 payload is also not immune... some windows carry a credible risk of catastrophic outcomes including potential loss of vehicle or loss of satellite."*[2] The advisory's own published partition gave its called catastrophic bin a 0.12 prior against 0.42 combined for the clean outcomes — a deliberate, dated minority position against a tracking consensus (NextSpaceflight, RocketLaunch.Live, the NSF Forum) that treated April 19 as uniformly cleared.

New Glenn had lifted off at 07:25 — inside Window 2.

It was the sort of fact that made juniors giddy and made Elif Kaya, who had underwritten space risk for eleven years, deliberately slow down. She had spent the weekend reading the rest of the record, including the advisory sealed three days earlier for the *previous* attempt — a day-wide no-go for April 18 with an eleven-window ranking that was never tested, because the operator stood down on its own criteria and the day came off the board. The record had graded its own untested ranking **NEAR, not HIT**, on exactly the grounds a skeptic would insist on: *no launch attempt was made in any named window.*[3] A record that marked its own near-misses down, she noted, was doing her grading discipline for her.

THE TWO READINGS

Whitaker's reading — the engineer's. Dr. Sam Whitaker, Aldgate's consulting propulsion engineer, had written his half of the memo in one sitting, and it was unanswerable on its own terms. *No known aerospace mechanism ties a rocket's minute-of-liftoff to the thrust of a second-stage burn that happens*

downrange, later, in vacuum. The upper stage would have underperformed at 07:18 or at 09:21. Window 2 "fired" because the T-3:57 hold and the forty-minute recycle happened to land the count inside it — a scheduling coincidence riding a coarser call. And the record *itself said so*: its own published analysis conceded the sub-window hit was not causal and priced only the day-level read as load-bearing.[4] As for the previous day's eleven-window ranking: never exercised, graded NEAR by its own author, zero information about whether intra-day orderings carry signal. Conclusion: for any desk whose decisions run through physical mechanism — a launch operator choosing slots, an engineer clearing hardware — this input class offered nothing an engineering review didn't, and adopting it would be theater. *For our purposes: decline.*

Kaya's reading — the underwriter's. Kaya agreed with every sentence of the physics and had signed nothing that disputed it. Her half of the memo began: *"Sam is right, and it doesn't matter, because our desk doesn't price mechanisms. It prices days."* Heritage-based launch rating — the industry's instrument — prices a vehicle by hardware lineage and flight count. The same booster earns the same rate every day it flies; the model is blind *by construction* to anything that varies launch-to-launch on identical hardware. A dated, falsifiable, day-level tail-risk read, published before the fact by a source that retained its misses, was native to exactly that blind spot. And the actions it supported on April 18 had been desk-level and same-afternoon: raise the attachment on this one flight; write a window-aware condition into the binder; require an independent reflight-verification step as a condition of binding. Hours of work, inside authority Aldgate already held, none of it pretending to fix an engine. The coarse read — *treat April 19 as elevated tail risk* — was real, was public, was a 0.12-prior minority position, and had landed. What she declined to adopt was precisely the part that looked most impressive: the nine-minute window. *For our purposes: adopt the day-level class, discard the sub-window dressing, and log every read against outcomes so we learn whether the class carries signal at our grain.*

The managing director read both drafts Thursday night and sent one line back: *"You've each told me what the input is worth at your desk. The memo I need says what Aldgate does — one rule, written down, that both of you will sign."*

EXHIBITS

Exhibit 1 — The mission record (public). April 14→16→18→19: T-0 moved four times. April 15: sealed day-wide no-go for April 18 with eleven-window ranking; hours later, operator stands down hot-fire on own criteria; April 18 comes off the board (ranking never tested; record grades its own sub-claims NEAR). April 19: T-3:57 hold, 40-minute recycle, liftoff 07:25 EDT inside flagged Window 2; clean booster landing; GS2 second-burn thrust shortfall; payload at ~154 × 494 km (McDowell tracking), declared total loss, insured value $30M; FAA grounding, later cleared. Sources: Spaceflight Now (Apr 20, 2026), Space Intel Report (Apr 19, 2026), SatNews (Apr 30, 2026), public tracking.

Exhibit 2 — The sealed input. Advisory of April 18, 2026 (LA-022), sealed via YouTube upload + X thread before the event: day no-go; ten windows; Window 2 (07:18–07:27 EDT) "high risk of catastrophic outcome"; five named anomaly categories including potential loss of satellite; published outcome-space partition (called bin prior 0.12 vs. 0.42 combined clean). Companion advisory LA-021 (April 15, for the scrubbed April 18 attempt) with its self-assigned NEAR grades. Full texts, grading, and the record's own coincidence concession at jyotishintelligence.com/intel/advisory/LA-022 and /LA-021. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The rating instrument (industry-standard, simplified). Heritage-based launch rating: rate = f(vehicle lineage, cumulative flights, payload class). Invariant across launch days by construction. The gap: any risk factor that varies launch-to-launch on identical hardware is unpriced.

Exhibit 4 — The two memo drafts (composite teaching device). Whitaker: mechanism-null, sub-window coincidental (per the record's own concession), LA-021 ranking untested → decline. Kaya: desk prices days not mechanisms; heritage rating blind by construction; day-level read actionable same-afternoon inside existing authority → adopt coarse class, discard fine grain, log and grade everything.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Decline (Whitaker)	B · Adopt the day-level class, graded	C · Adopt with sub-windows
What it adds	engineering purity	a lane for the one input class heritage rating can't see	impressive-looking precision

What it forecloses	the blind spot stays	nothing — every read logged and graded	the grain discipline
Cost if the class is noise	0	a log that proves it, in time	coincidences institutionalized
Cost / value if day-structure is real	binding blind through flagged days	priced attachments, graded learning	credit claimed for scheduling luck
Speed	—	same-afternoon desk actions	same
Authority required	none	desk's own	desk's own (unfortunately)
Reversibility	n/a	full	sheets are hard to un-teach
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/two-readers
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — Two Readers, One Input

Rails: the scheduling-coincidence concession is the case's center of gravity and comes from the record itself — a telling in which "the astrologer called the nine-minute window" survives has inverted the case; LA-021's NEAR grade must stay visible (the ranking instrument was never tested, anywhere); space is hard, nobody in the chain is a villain, and the physics half of the room must win its argument *completely* for the case to work.

LEARNING OBJECTIVES

(1) Information is reader-relative: an input's decision value is a function of the *decisions a desk actually makes*, not of the input's accuracy — the same document can be worthless at the engineering desk and priceable at the underwriting desk, with both readings correct; (2) grain discipline: how to consume a warning at the resolution its evidence supports (the day) while explicitly discarding the resolution it doesn't (the nine-minute window), even when the finer grain "hit"; (3) being right can add nothing: a striking confirmation that rides a coincidence carries zero incremental information, and a record honest enough to say so about its own hit is exhibiting the only property that makes any of its other grades worth reading; (4) instrument blindness as the locus of value: warnings are worth most exactly where the incumbent pricing model is blind by construction.

ASSIGNMENT QUESTIONS

1. Write the one rule both authors will sign. It must tell a future Aldgate underwriter what to do with the next sealed day-read, at what grain, with what grading obligation.
2. Whitaker is completely right about the physics. Construct the strongest version of his "decline" — then find the premise about *what Aldgate's decisions are* that it quietly assumes.
3. The liftoff landed inside Window 2. A junior proposes citing this in marketing. Write the two-paragraph reply that explains why the desk will not, using only the record's own published concessions.
4. LA-021's ranking graded NEAR because the day was scrubbed before any window opened. What experiment — over how many missions — would distinguish an intra-day ordering that carries signal from one that doesn't? Who has the data, and why hasn't anyone run it?

5. Generalize Exhibit 3: name a pricing instrument in your own industry that is blind by construction to a variable class, and describe what a dated, falsifiable read of that class would be worth at your desk — independent of its source.
6. **The machine staff-officer.** The agency's AI pricing model will happily ingest day-level reads as features — and with enough data it will also rediscover Window 2 as 'signal,' because scheduling coincidences correlate with holds and recycles. Encode Whitaker's physics and Kaya's grain discipline as an actual feature-governance rule: which resolutions of this input class may enter the model, what evidence threshold promotes a finer grain, and who reviews it?

THE BOARD PLAN (90 MIN)

Open with the raw sequence (10) and let the room gasp at Window 2 — the case needs the gasp on the record so it can be dismantled. Whitaker's argument (15): build it to full strength; the physics must win. Then the turn (15): Kaya's "right, and it doesn't matter" — walk Exhibit 3 until the class sees that the two readers are answering *different questions*, and that information economics, not physics, is the case's subject. The grain discipline (15): Q3, written; this is where objective 3 lands, and the record's own concession does the work. The rule (20): Q1 drafted and defended; press until it handles the next input, not this one. Close (15) on Q5 — transfer to the students' own industries, which is the whole point of the case.

WHAT HAPPENED, AND THE GRADE

There is no hidden outcome to reveal — the case's events are complete on its first page, which is itself the pedagogy: this is the one case in the volume where the decision under study is not what to do before an event but what to *learn* from one. The record grades LA-022's day-level call a HIT at its published 0.12 prior (~3.1 bits against the tracking consensus), grades the companion LA-021 ranking NEAR because reality removed the test, and concedes in its own text that the Window-2 confirmation was a scheduling coincidence riding the coarser read. The FAA investigation ran its course and New Glenn was later cleared. Whether Aldgate-class desks should adopt day-level reads is a live question the record itself answers only with an invitation to keep score — which is, precisely, the rule the best student memos converge on.

EPILOGUE

The deepest sentence available in this case belongs to neither protagonist but to the grading ledger: the same record that reports the hit reports, a few lines away, its own nine misses and its own NEAR — kept at full weight, recomputable by a stranger. Strip that away and the case collapses into an anecdote about a lucky window. Keep it, and the case becomes what this volume is for: a demonstration that the question "was it right?" is almost never the operative one. The operative questions are the two desks' —

does it reach a decision I actually make? and *at what grain does the evidence hold?* — and they have different answers in the same room, on the same morning, about the same nine minutes.

CASE TWELVE

The Founder Window: A Meeting on the Calendar

A teaching case in concentration risk, key-person exposure, and the deadline posture — converting an open-ended worry into dated work

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the sealed input in this case is, by the record's own choice, **ungraded**: no probability was staked at seal, and the ledger declines to backfill one. Every public event, figure, and statement in the case is real and cited; the named public figures appear only through their public record, and the case depicts no real agency's internal deliberations and no villains. The **protagonist and her office are composite characters** — "Dana Reyes" and the "Portfolio Assurance Cell" are not real; their portfolio and meetings are teaching devices, and no outside-world fact they discuss is invented. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

THURSDAY, MARCH 20, 2025 — 14:00, WASHINGTON

Dana Reyes ran supplier-concentration analysis in a small federal acquisition-support cell — the desk whose job was to know, before anyone asked, what happens to a mission when a single vendor's circumstances change faster than the acquisition system can react. Her portfolio's most concentrated name was not a secret; it was arguably the most successful company in American aerospace. SpaceX carried sole-source crew transport to the International Space Station and a reported federal contract book in the tens of billions, and it carried them well. Every instrument on Reyes's dashboard — financial health, delivery performance, schedule — read green, because every instrument on her dashboard was built to detect a *failing* supplier, and SpaceX was the opposite of failing.

Which was exactly the blind spot the strange item in her inbox pointed at.

A colleague who tracked open-source chatter had forwarded, with the note *"filing this under 'probably nothing, costs nothing,'"* a public, timestamped post of **March 19** from a source maintaining a sealed and publicly graded forecast record. It read, in the words preserved at the source: *"Keep an eye on June 2025, especially between June 8 and 10"* for *"impending misfortunes with SpaceX"* — explicitly distinguished from Tesla, with the founder personally invoked.[1] The source's method, disclosed in its own materials, was Vedic astrology.

Reyes made herself inventory the sealed words at their own level, no more: one named company. One named founder. One three-day window, seventy-eight days out. One word of direction — *misfortunes* — with no mechanism, no probability, no government angle, no target. And she inventoried the context, which was what made the statement strange rather than merely thin: in March 2025, Elon Musk and the President were close allies, fresh off the DOGE effort; Musk was at the peak of his public influence; SpaceX was private, ascendant, with no tradeable instrument on it and no public analysis carrying any rupture of any kind. A June misfortune window for SpaceX ran against essentially all available information — which meant it was either noise, or the only statement in her intake that wasn't already priced into her dashboard.

THE BLIND SPOT, NAMED

What the flag surfaced, whether or not it deserved belief, was a category error in her own toolkit. Vendor concentration asks: *what if the company fails?* Her instruments answered that daily. Founder concentration asks a different question: *what if the company is fine — and the founder is the event?* For a founder-led sole-source supplier, the portfolio was, in effect, silently long one person's disposition, and no clause, dashboard, or program review recognized that position, because nothing in the acquisition system had ever been asked to. The sealed flag, by naming the man and not just the company, was inviting exactly that reading.

THE THREE POSTURES

She priced what the flag could buy, holding every action to what the sealed words supported.

Package A — Monitor. File it with the colleague's note. Cost: zero. In the bad world: whatever work the window demanded would be performed inside a news cycle instead of across eleven weeks.

Package B — Quiet readiness, on a deadline. Five items, all desk-level, all inside existing authority, all schedulable across April and May precisely *because* the window was dated. *One:* a continuity-disruption map — what mechanically happens if this contractor becomes unable or unwilling to perform, for any reason; which clauses speak; what the government actually holds in data rights, tooling, and trained alternates. Counsel hours. *Two:* an alternate-provider status refresh at liaison pace — the genuine readiness of every alternate on the same contract lines, not the brochure version. One to two weeks. *Three:* a key-person clause audit — where, across the affected contract set, any instrument recognized the founder as a risk locus at all. Days. *Four:* a pre-drafted holding position — continuity assured, options under review — written in April, when it cost hours, rather than sought in June, when it would be unavailable at any price on a short fuse. *Five:* the cheapest item in this book — **a calendar entry**: a standing agenda item for the first week of June, *review the continuity annex before the flagged window*. If the window passed empty, the total cost of item five was a meeting.

Package C — Elevated posture. Escalate the flag itself: brief leadership, open formal alternate-sourcing analysis, treat the window as tasking. Cost: her credibility, the cell's alarm budget, and a public signal the flag's evidentiary weight could not carry. Struck through.

The objection came from her own deputy, in the form the acquisition world respects: *"We do not task counsel off an astrologer. If this leaks, the sentence 'the government audited SpaceX continuity because of a horoscope' ends this office."* Reyes's draft answer, which she was still testing when the meeting ended: every item in Package B was work the desk should arguably have had standing anyway for its most concentrated name; the flag's only real contribution was a *deadline* — a date that converted "we should someday look at founder risk" from an unfunded intention into a bounded, schedulable, June-first deliverable. Deadlines, not beliefs, are what move desk work. The question was whether she could write that rationale into the file in a way that would read as prudence in every possible world — including the one where June came and went without incident, and the one where the file someday surfaced in an oversight hearing.

The April work-plan went to her director Monday.

EXHIBITS

Exhibit 1 — The consensus, March 19, 2025. Musk–President alliance at its public peak post-DOGE; SpaceX private, ascendant; sole-source crew transport; reported federal exposure ~$22B across programs; no public analysis carrying any rupture. Every conventional supplier-risk instrument green.

Exhibit 2 — The sealed input. Public post of March 19, 2025 (x.com/vijayjyotishusa/status/1902444815897530872): the June 8–10 window, "impending misfortunes with SpaceX," founder personally invoked, no mechanism or probability stated. Advisory documentation at jyotishintelligence.com/intel/advisory/IA-SX-001 — including the record's own **refusal to grade the entry** (no probability staked at seal; none backfilled after the fact, although the outcome would have flattered the aggregate). Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The instrument gap (composite teaching device). Vendor-concentration toolkit: financial-health monitoring, delivery/schedule performance, single-award ratios — all keyed to supplier *capability*. Unkeyed anywhere: founder disposition, change-of-control, key-person triggers, communication protocol with the contractor's governance layer.

Exhibit 4 — What the window carried (public record, for the teaching note). June 5–7, 2025: the alliance ruptures in public; the President floats cancelling SpaceX's federal contracts; Musk posts — and deletes within hours — that SpaceX will "begin decommissioning its Dragon spacecraft immediately"; NASA and the Pentagon reported sounding out competitors mid-crisis; Tesla sheds a record $152B in one session; Musk personally −$33.9B; a longtime investor's same-day estimate puts SpaceX's private valuation "probably cut in half." Walk-back within a day. Twelve months later, June 11, 2026: SpaceX goes public in a record ~$75B raise at $135/share. Sources: Reuters (Jun 5, 2025), Washington Post (Jun 7, 2025), CNBC (Jun 5, 2025), Fortune (Jun 6, 2025), NPR (Jun 11, 2026).

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · Deadline package	C · Escalate the flag

What it adds	nothing	continuity map + alternates review + clause audit + comms pre-draft + a meeting on the calendar	leadership attention, now
What it forecloses	desk-tempo conversion	counsel hours, one agenda slot	the office's credibility on one unverifiable flag
Cost if June is quiet	0	work the portfolio arguably owed its most concentrated name anyway	the leak scenario
Cost / value if the window lands	the ten days, in public, at crisis tempo	activation instead of drafting	positioned — if believed
Speed	—	weeks, off-cycle	principals' calendar
Authority required	none	her own + counsel	leadership
Reversibility	n/a	full	poor
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading — the record itself refused to)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/founder-window
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Founder Window

Rails: the entry is **ungraded by the record's own choice** and the note must teach that refusal as seriously as any grade; the measurable financial evidence runs through Tesla, a correlated proxy — SpaceX-specific damage was never directly measured; the cluster broke June 5–7 against a named June 8–10 window (early, inside the week; the record claims no sub-window precision); "misfortunes" is semantically elastic; the political principals appear only through the public record and the classroom takes no side in their dispute.

LEARNING OBJECTIVES

(1) The deadline posture: a dated window's chief product is not prediction but *scheduling* — it converts an open-ended, chronically deferrable risk question into bounded work with a due date, and deadlines are what actually move institutional desks; (2) founder concentration as a distinct risk class from vendor concentration, invisible to capability-keyed instruments by construction; (3) tempo economics: the same work performed at desk tempo in April versus crisis tempo inside a news cycle — the case's entire counterfactual is conversion of tempo, never prevention; (4) how to weight an input that stakes no probability — and what a record's refusal to score its own flattering outcome teaches about the record.

ASSIGNMENT QUESTIONS

1. Write the rationale paragraph for Reyes's April work-plan — defensible in the empty-window world, the crisis world, and the oversight-hearing world simultaneously. May it mention the flag's source? Must it?
2. The deputy's leak scenario is real. Price it against the crisis-tempo scenario. Which asymmetry governs, and what does that say about *how* unconventional inputs should be filed, cited, and laundered — or pointedly not laundered?
3. Package B item five costs a meeting. Steelman the case that even a meeting is too much — that intake rules admitting this input class admit an unbounded stream of noise. Then write the intake rule that bounds the stream without excluding the class.
4. "The rupture was foreseeable in kind to any newspaper reader — the flag only supplied a date." Grant it. For a *deadline* posture, is a date not precisely the scarce input? Where does this argument leave inputs that supply direction but no date?

5. The record refused to grade this entry because no probability was staked, though the outcome would have improved its aggregate. What does that refusal predict about the record's other entries — and what analogous discipline should Reyes's own office adopt for its warnings?
6. **The machine staff-officer.** An AI concentration-risk dashboard now scores founder disposition continuously — sentiment, litigation, political proximity — and flags 'windows' weekly. Continuous scoring without deadlines reproduces Package A at higher cost: everything is always slightly urgent. Design the protocol that converts machine-flagged windows into calendar entries with owners and expiries — and state the threshold below which flags must not interrupt anyone.

THE BOARD PLAN (75–90 MIN)

Open on Exhibit 3 (15): build the founder-vs-vendor distinction until the class can name why every dashboard was green and why green was the wrong color to trust — no forecast on the board yet. Introduce Exhibit 2 at its own level (15): the four sealed elements, the staked-nothing problem, Q4's objection surfaced early. The three postures (15), with the deputy's leak scenario given full dramatic weight — Q2. The work-plan rationale (20): written commitment; the strongest drafts discover that Package B needs the flag only as a *date*, which is objective 1 achieved from inside the exercise. Reveal Exhibit 4 (10): the room will want vindication; walk them instead through the concessions (proxy evidence, early window, elastic word, ungraded entry). Close (15) on Q5 — the refusal to score as the deepest fact in the case.

WHAT HAPPENED, AND THE GRADE

There is a happened and deliberately no grade. On June 5–7, 2025 — the leading edge of the named week, two to three days ahead of the sealed window — the alliance ruptured in public, exactly along the axis the portfolio was silently long: for one evening, the continuity of United States crew transport to the ISS turned on one founder's disposition, expressed through a post no contracting officer or clause could intercept between composition and publication. The contingency work of Package B was ultimately performed — by the government, at crisis tempo, partly in public, inside the news cycle. The threat was walked back within a day; no service was interrupted; the following June the same company priced the largest stock-market debut on record. The record holds IA-SX-001 **ungraded**: the flag staked no probability, so the ledger — by its own published rule — refuses it a score it would have benefited from. The teaching synthesis stands on the asymmetry that survives every concession: a dated, falsifiable, adverse flag on the portfolio's most concentrated name existed in public seventy-eight days before that supplier's worst week in memory; the hedge it invited cost counsel hours and a calendar entry; and the work it invited was done anyway — later, faster, and in front of cameras.

EPILOGUE

This is the last case in the volume, and it closes where the first one opened: with the cheapest artifact on the board. A meeting on the calendar — *review the annex before the flagged window* — is the deadline posture reduced to its atom. It requires no belief, survives every skeptic in the room, and converts the one thing an unverifiable warning indisputably contains, a date, into the one thing a desk indisputably runs on, a due date. Students who leave remembering nothing else should leave with the closing question every case in this book has been asking in a different uniform: *what is the cheapest thing I could put on the calendar today that future-me, standing in the one bad week, would call the best purchase this office ever made?*

CASE THIRTEEN

The Exit Polls: A Ceiling Under the Consensus

A teaching case in political risk, scenario planning, and building the branch the room refuses to price

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the sealed reading in this case carries **three graded MISSES and five PARTIALs** among its sub-forecasts, all preserved and all cited. Every election figure, poll number, and market move is real and cited. The **protagonist and her firm are composite characters** — "Anita Rao" and "Vantara Industries" are not real; their regulatory agenda and meetings are teaching devices, and no outside-world fact they discuss is invented. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works, and no position on any Indian political party.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

THURSDAY, AUGUST 17, 2023 — 11:00, MUMBAI

Anita Rao ran government affairs for Vantara Industries — a large diversified firm with a regulatory agenda that lived or died in Delhi: two sector reforms in draft, one licensing question pending, and a five-year plan whose assumptions all began with the same sentence. *The government's direction is known.* Nine months out from India's 2024 general election, that sentence was not controversial anywhere she could hear. The ruling alliance was campaigning toward a supermajority — the slogan would soon be "Abki Baar 400 Paar," four hundred seats or more — and every institutional analyst her firm retained treated a third term with an enlarged single-party majority as the planning baseline. Her board deck said so. Her competitors' decks, she assumed, said the same.

The item her analyst flagged that morning cut under the consensus rather than against it. A public forecast thread of **August 10**, timestamped on the platform, from a source with a sealed and publicly graded record, predicted the consensus headline — Modi wins a third term — and then put a ceiling

where everyone else had a floor: the alliance's 372/400+ target was, in the words preserved at the source, *"just a dream,"* and the party *"would be well advised to set realistic goals of 300+ and optimize the strategy for retaining the government at all costs."*[1] The thread carried roughly thirty-two falsifiable sub-forecasts — state calls, a funding-scandal axis, security windows. The source's method, disclosed in its own materials, was Vedic astrology.

Rao noticed what the ceiling implied for *her* desk specifically, because it was structural, not electoral. A single-party supermajority needs no partners: her playbook under that outcome was the playbook she already had. But a government *retained at all costs* — a coalition-dependent alliance clearing 300 with allies — was a different regulatory universe: every contested bill priced in ally votes, sector asks riding on coalition partners' state-level agendas, reform timelines stretched by bargaining. Two futures, two playbooks — and her firm had built exactly one.

THE THREE POSTURES

Package A — Monitor. File the thread. Keep the single-branch plan. Cost now: zero. Cost if the ceiling landed: building an ally-leverage map in the compressed weeks after the result, while every competitor scrambled at the same doors.

Package B — The quiet contingency. Two items, both inside her own signature. *One:* an ally-leverage map, built at the level the sealed reading itself supported — not naming winners, but mapping in general terms which kinds of allied parties plausibly hold marginal leverage under a coalition outcome, and what coalition partners historically ask for in Indian federal politics (state fiscal asks, ministry allocations, sector carve-outs), cross-indexed to Vantara's agenda. Analyst-weeks, in the August lull before campaign season claimed everyone's calendar. *Two:* a messaging hedge — holding the firm's public affairs and sponsorship language to outcome-agnostic phrasing rather than supermajority triumphalism, at zero cost beyond discipline.

Package C — Pre-emptive repositioning. Rebuild the five-year plan around the coalition branch; brief the board that the supermajority was doubtful. Cost: her credibility staked on a single anonymous-to-the-board thread against a unanimous professional field, and a public posture the firm could not walk back gracefully.

Her deputy's objection was the room's: *"Every pollster, every bank note, every retained consultant says 350-plus. You want desk-weeks against all of them on an astrologer's thread?"* Rao's draft reply was that the thread was not the argument — the *asymmetry* was: the coalition branch cost an afternoon to open and analyst-weeks to furnish, in a quiet season; it cost nothing to discard if June proved the consensus right; and if June proved the consensus wrong, the map could not be bought later at any price that mattered, because the post-result bargaining would already have started. What the thread contributed was not certainty but a dated, falsifiable, specific ceiling — the only number in her intake that diverged from the field — and, she wrote carefully, *plain anti-incumbency reasoning after ten years of single-party rule pointed the same direction anyway.* The question for Monday's planning session was which packages to commission.

EXHIBITS

Exhibit 1 — The consensus, August 2023 → June 2024. Institutional baseline through the campaign: enlarged BJP majority. Three days before counting (June 1, 2024): exit-poll field converged on NDA 350–400+; Axis My India 361–401; Zee's promoted AI-anchored exit poll ("Zeenia") 305–315. Every public number, human or machine, above the eventual result.

Exhibit 2 — The sealed input. Public thread of August 10, 2023 (root x.com/vijayjyotish/status/1689539625004482560, T3–T19/55): third term YES; 372/400+ "just a dream"; "realistic goals of 300+... retaining the government at all costs"; ~32 falsifiable sub-forecasts including state calls and a funding-scandal axis. Advisory documentation, grading, and the full honesty ladder — including the West Bengal MISS and the seven other adverse grades — at jyotishintelligence.com/intel/advisory/IA-IN24-001. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — What the two branches cost (composite teaching device). Ally-leverage map: analyst-weeks in August 2023; the same map in June 2024 = compressed, contested, late. Messaging hedge: zero marginal cost. Board repositioning: credibility-priced, hard to reverse.

Exhibit 4 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · Quiet contingency	C · Repositioning

What it adds	nothing new	second playbook, seasoned early	board alignment to the ceiling
What it forecloses	the early map	nothing	graceful walk-back
Cost if nothing happens	0	analyst-weeks	credibility spent against the field
Cost / value if it happens	late scramble, contested window	bargaining-ready on day one	fully positioned — if believed
Speed	—	weeks	months + board cycle
Authority required	none	GA head's own signature	board
Reversibility	n/a	full (discard the memo)	poor
Recommendation & rationale		*(yours)*	
In hindsight		*(filled at grading)*	

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/exit-polls
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Exit Polls

Rails: the case takes no side in Indian politics and grades no party; the sealed reading's eight adverse grades (three MISS — West Bengal cleanly among them — five PARTIAL) must reach the board at full weight; the honest alternative route must too: plain anti-incumbency base-rate reasoning could have produced the same *directional* skepticism without any forecast — what it could not produce is the specific ceiling band, or a dated artifact to plan against.

LEARNING OBJECTIVES

(1) Ceiling-versus-floor reading: the highest-value divergence is often not "the consensus is wrong" but "the consensus is unbounded" — a ceiling under an agreed outcome changes the *playbook*, not the winner; (2) contingency branches as options priced by asymmetry: cheap to open in the lull, unbuyable in the scramble; (3) unanimity as information: when every instrument — including the AI one — sits on the same side, the marginal value of one dated dissent is at its maximum precisely because the room's hedges are at their minimum; (4) grading discipline on a mixed record: how to weight an input whose own ledger shows 24 hits, 5 partials, and 3 misses across its sub-forecasts, none hidden.

ASSIGNMENT QUESTIONS

1. Write Rao's Monday commissioning memo. Which packages, on what stated basis — and does the thread appear in it by name?
2. Grant the deputy's objection fully: the entire professional field, with ten more months of data, disagreed. Construct the strongest case for Package A — then locate the premise about *cost symmetry* it silently assumes.
3. The base-rate route (anti-incumbency after a decade) reaches the same direction free. What, precisely, did the sealed ceiling add that the base rate did not? Price each component (the band, the date, the artifact).
4. The same thread carried a clean West Bengal MISS. You are Rao in July 2024, re-evaluating the source for future intake. Write the two-paragraph evaluation that neither over-credits the ceiling hit nor waves away the misses.
5. Complete Exhibit 4: sign a recommendation and defend the C column's honesty — if C looks absurd, is the menu rigged?

THE BOARD PLAN (75 MIN)

Open on the two playbooks (15): make the class own that a supermajority and a coalition are different *regulatory* worlds before any forecast appears — the case is about branch-building, not election-calling. The consensus wall (10): Exhibit 1, including the AI exit poll — students should feel the unanimity. Introduce Exhibit 2 with its full honesty ladder (15): run Q3 — the base-rate objection is the session's intellectual center and should be given its full force. Price the packages (15): Q2's asymmetry argument, drawn as a two-world table on the board. Commit (15): Q1 memo and the Exhibit 4 signature. Close (5) on Q4 — the mixed-record weighting question that every real intake function eventually faces.

WHAT HAPPENED, AND THE GRADE

On June 4, 2024 — 299 days after the seal — the count came in at **NDA 293, BJP 240**: the party's first loss of a single-party majority since 2014, forcing genuine coalition government. The exit-poll field had converged three days earlier on 350–400+; every public number, human or machine, sat above the result; the Sensex fell roughly 5.7% on counting day — close to $386 billion intraday — as the professional consensus repriced. The record grades the ceiling call a HIT within a reading it scores honestly: of ~32 sub-forecasts, three MISS (West Bengal cleanly; a named security window that produced no event; an unverifiable intelligence claim) and five PARTIAL, all retained at full weight — and its published derivation prices the joint call at roughly 1-in-4,400 before the fact (~12 bits) against the specific consensus of August 2023. For Rao's composite desk, the postures priced exactly as the asymmetry predicted: a firm holding the ally-leverage map entered the post-result bargaining with its homework done; the field built the same map afterward, in the scramble, at bargaining prices.

EPILOGUE

The AI exit poll deserves its own coda: Zeenia beat every human pollster and still finished above the result — the machine had compressed the field's error, not escaped its direction. The teaching point is the volume's recurring one, wearing electoral clothes: the scarce commodity was never a better point estimate. It was a dated, falsifiable *branch* — held open cheaply, early, against unanimity — and the desk discipline to furnish it before the world announced which branch was real.

CASE FOURTEEN

The Substitute Port: Enforcement Before the Reroute

A teaching case in sanctions policy, evidentiary tempo, and pre-packaging the measures a campaign will eventually demand

> **AUTHORSHIP DISCLOSURE.** This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the sealed claim's **maximal end-state never occurred**: the bridge was struck, repaired, and kept operating, a shortfall the record's own grading notes carry and this case keeps in view throughout. Every strike, date, and public statement is real and cited. The **protagonist and her unit are composite characters** — "Marta Keller" and the "Coalition Enforcement Cell" are not real; no actual ministry's deliberations are depicted, and no outside-world fact they discuss is invented. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works.
>
> **METHOD DISCLOSURE.** The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

MONDAY, SEPTEMBER 11, 2023 — 09:30, A G7 CAPITAL

Marta Keller ran a small desk inside a coalition enforcement cell — the working level of the G7 oil-price-cap regime, where designation dossiers were drafted, maritime screening analytics were tasked, and agenda items for the coalition's coordination calls were fought over one slot at a time. Her professional life ran on two clocks that nobody outside the building believed: **evidence assembly for a designation takes months**, because a package that cannot survive litigation is worse than no package; and **coalition coordination is slower still**, because measures that land jointly bite and measures that land alone leak.

The input her intake analyst had flagged the previous week was, on its face, someone else's problem — a military matter, not a financial one. A public, timestamped post of **September 4**, from a source with a sealed and publicly graded forecast record, stated: *"There is a potential for an all-out attack on the*

Kerch Strait bridge in Crimea, with the goal of affecting damage beyond the point of repair."[1] Two elements fixed in one sentence: a campaign character — all-out, sustained, not episodic — and a strategic intent, permanent disabling of the land-bridge into occupied Crimea. The source's method, disclosed in its own materials, was Vedic astrology. The bridge had been struck twice already — the October 2022 truck bomb, the July 2023 surface-drone attack — but Ukrainian doctrine was publicly unsettled between symbolic and operational targeting, and no naval-analysis institution had a dated, public campaign-with-disabling-intent call on the record.

What made it her problem was the shadow the sentence cast. Keller's desk did not care about bridges; it cared about **where Russian maritime activity concentrates**. A campaign that degraded the Kerch land-bridge and kept pressure on Sevastopol would predictably push Russian naval and commercial traffic toward whatever Black Sea port the campaign did not touch — a substitution pattern, forming over months, on exactly the AIS and port-call surfaces her analysts already watched. The sealed sentence named no port and no sanctions consequence; the substitution logic was her own inference, and she wrote it down as such. But if the inference held, then the enforcement work the reroute would eventually demand — screening packages, designation dossiers on the facilitation networks that would service the substitute node, coalition alignment — had lead times of *months*. Which meant the only day cheap enough to start was a day like today, when nothing had happened yet.

THE THREE PACKAGES

Package A — Monitor. Track the campaign in the weekly read. Commission nothing. Cost now: zero. Cost if the campaign landed as sealed: the desk starts its evidentiary clock on the day the reroute is publicly confirmed — months behind the pattern, chasing a substitution channel that has already hardened.

Package B — Quiet pre-packaging. Three lines, all inside existing authority, none requiring legislation, ministerial risk, or public commitment. *One:* a **port-substitution screening package** — task AIS and port-call analytics on Black Sea tonnage shifting load patterns away from Kerch-adjacent flows toward Russia's other Black Sea ports; flag insurance-attestation review for vessels changing load ports. Analyst-weeks. *Two:* **contingency designation dossiers** — open evidence files now on the facilitation-network classes that would service a substitute export node; the long pole in any designation is evidence assembly, and the file's own arithmetic said the lead time roughly

equaled the pole. *Three:***coalition sequencing** — one agenda item, this autumn: a port-substitution contingency annex, so that if the pattern formed, the coalition's slowest clock had already been started. Total cost: staff time and one agenda slot.

Package C — Declared posture. Take the campaign call to principals; propose public anticipatory measures. Cost: ministerial exposure staked on a single-source, unconfirmed campaign forecast, and a public signal that would itself teach the adaptation it was meant to catch.

The objection at Thursday's tasking review came from her senior analyst, and it was good: *"After two strikes on that bridge, 'more strikes coming' is momentum, not information. And 'beyond the point of repair' is a maximal claim nobody serious would underwrite."* Keller conceded half of it on the spot — the maximal end-state was exactly the kind of clause a graded record would eventually have to carry as a shortfall. What the momentum account did not supply, she argued, was the *campaign framing* — sustained, all-out, aimed at permanence — and the campaign framing was the only part her desk could use, because a durable reroute is driven by sustained campaigns, not episodic strikes. Episodic strikes produce headlines; campaigns produce port substitution. The question on the tasking sheet was whether that distinction was worth three lines of staff work.

EXHIBITS

Exhibit 1 — The strike history and the doctrine gap, as of September 4, 2023. October 2022: truck-bomb attack on the bridge; July 2023: surface-drone attack. Ukrainian targeting doctrine publicly unsettled between symbolic and operational; naval-analysis institutions (order-of-battle tracking, capability analyses) carrying no dated campaign-intent call.

Exhibit 2 — The sealed input. Public post of September 4, 2023 (x.com/vijayjyotish/status/1698760258095571024), quoted verbatim in the case text. Advisory documentation and grading — including the record's own caveat that the beyond-repair end-state has NOT occurred — at jyotishintelligence.com/intel/advisory/IA-RU-019. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The two clocks (composite teaching device). Designation evidence assembly: months (litigation-grade quality is the binding constraint; time pressure is its enemy). Coalition coordination: the slowest clock in the

system — measures land jointly or leak. Screening package: analyst-weeks. Insurer advisory: days.

Exhibit 4 — What subsequently happened (public record, for the teaching note). Apr 3, 2024 (D+212): Ukraine's defence-intelligence chief, on the record: the bridge "will be disabled in the first half of 2024." May 17, 2024: long-range ATACMS delivery confirmed. May 29, 2024 (D+268): ten-missile ATACMS salvo against air defences protecting the bridge approach. May 30, 2024: ATACMS strike destroys two rail ferries at the Kerch crossing. Jun 3, 2025: SBU underwater strike on the bridge supports; traffic resumed same day. Apr 18, 2024: UK Defence Intelligence publicly describes the Russian Black Sea Fleet relocating key operations to Novorossiysk — the principal commercial export node the campaign did not touch.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · Quiet pre-packaging	C · Declared posture
What it adds	nothing new	screening + dossiers + coalition clock started	public deterrent signal
What it forecloses	the evidentiary head start	nothing	quiet optionality; teaches adaptation
Cost if nothing happens	0	staff time, one agenda slot	ministerial credibility
Cost / value if it happens	months behind a hardened channel	packages ready as the reroute confirms	positioned — if the single source held
Speed	—	days / weeks / one agenda cycle	principals' calendar
Authority required	none	desk + unit head	ministers
Reversibility	n/a	full (files close quietly)	poor
Recommendation & rationale		*(yours)*	
In hindsight		*(filled at grading)*	

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/substitute-port
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Substitute Port

Rails: the sealed maximal end-state **did not occur** — the bridge was repaired and kept operating; the record carries that shortfall in its own grading notes and the classroom must too; the substitution inference (which port) was Keller's, never the seal's — the sealed-text line must stay bright; the momentum-extrapolation objection gets full force; this is an enforcement-posture case, and no protective or humanitarian claim is made anywhere in it.

LEARNING OBJECTIVES

(1) The package posture in government: when the binding constraints are evidentiary and coalitional clocks, a warning's value converts almost entirely into *start dates* — the cheapest thing a long-lead input buys is the right to begin assembling what the event will demand; (2) consequence-class reasoning: how to act on the *shadow* of a claim (substitution patterns) without exceeding what the sealed words support (no port named); (3) campaign versus episode: why sustained-intent framing, not another strike prediction, was the only enforcement-relevant content — durable reroutes are made by campaigns; (4) partial-fulfillment discipline: how to use an input whose direction confirmed and whose maximal clause did not, without rounding it either way.

ASSIGNMENT QUESTIONS

1. Write Keller's tasking-sheet justification for Package B — three sentences per line item, each pitched at what the sealed sentence itself supports.
2. Steelman the senior analyst completely: momentum plus a maximal clause nobody would underwrite. Now identify the one word in the sealed sentence the momentum account cannot generate, and what that word is worth on *this* desk.
3. Package C "teaches the adaptation it was meant to catch." Unpack that: when does announcing enforcement readiness reduce evasion, and when does it accelerate it?
4. The end-state clause failed; the campaign clause held; the reroute formed. Grade the input yourself, on the record's own terms — then decide what your grade changes about the next tasking decision.
5. Complete Exhibit 5 and sign it. Then write the "in hindsight" row twice — once for the world of Exhibit 4, once for a world where the campaign fizzled by December.
6. **The machine staff-officer.** An AI system with access to AIS feeds, sanctions lists, and the sealed input drafts Package B's screening package overnight and pre-fills Exhibit 5. Which cells do you accept, which do you audit, and which — exactly —

must remain human? Write the one-paragraph standing order that governs it.

THE BOARD PLAN (75–90 MIN)

Open on the two clocks (15): make evidence-assembly tempo and coalition tempo physical for the room — students without government experience consistently assume enforcement is fast. The sealed sentence, decomposed (10): campaign character vs. maximal intent, and the substitution shadow as *inference*, labeled as such on the board. The three packages (15). The momentum fight (15): Q2, adversarial pairs. Commit (15): Q1 and the Exhibit 5 signature. Reveal Exhibit 4 (10): run Q4 — the partial-fulfillment grading is this case's distinctive move; insist the class grade the clauses separately. Close (10) with Q6 — the machine can watch every port on earth; it cannot sign the designation.

WHAT HAPPENED, AND THE GRADE

The campaign framing held; the maximal clause did not. Ukraine's defence-intelligence chief publicly declared the disabling intent 212 days after the seal; ATACMS salvos struck the bridge's protective air defences and the Kerch rail ferries in late May 2024; a third major strike hit the bridge supports in June 2025 — and the bridge, each time, was repaired and kept partially operating, so *"beyond the point of repair"* stands unachieved, carried as an explicit caveat in the record's own notes. The enforcement-relevant consequence landed as the inference predicted: under cumulative pressure the Black Sea Fleet relocated key operations to Novorossiysk — publicly described by UK Defence Intelligence in April 2024 — concentrating Russian maritime activity on the principal export node the campaign never touched. A desk that had opened Package B in September 2023 would have been watching the substitution pattern *form* and holding designation files months into their evidentiary pole when confirmation arrived; a desk that waited started seven months later, against a channel that had spent those months hardening. Adaptation time is the adversary's cheapest resource, and reactive enforcement donates it.

EPILOGUE

This case's residue is a sentence Keller's trade knows and every other desk in this book keeps rediscovering: **lead time is the only input that shortens the slowest clock.** Nothing about the sealed sentence made the coalition faster, the evidence easier, or the litigation safer. It only moved the start date — which, on a desk where everything takes months, was the entire game.

CASE FIFTEEN

The Grind, Not the Gamble: An Estimate Meets a Failed Warning

A teaching case in intelligence analysis, indicator design, and what to do with a source whose testable half broke

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and this case is built on a sealed document whose **dated, falsifiable component did not occur**: the March–June 2024 attack it named never happened, the record says so in the first line of its own file, and no averted harm is claimed by the record, by this case, or by anyone in it. Every troop figure, talks round, and public statement is real and cited. The **protagonist and her staff are composite characters** — "Devika Menon" and the "Assessments Cell" are not real; no actual intelligence organization's deliberations are depicted, and no outside-world fact they discuss is invented. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works, and no position on India–China policy.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

MONDAY, NOVEMBER 6, 2023 — 10:00, A NATIONAL ASSESSMENTS CELL

Devika Menon owned a standing estimate — the document, revised quarterly, that answered one question for her principals: *will China move kinetically on the eastern Line of Actual Control?* Three years into the standoff, her indicator set was a mature piece of craft, and it was built the way warning sets are always built: around the event. Mobilization signatures, logistics buildup, forward staging, missile-force dispersal. If the People's Liberation Army came, her tripwires would see it coming.

The document her intake officer logged that week claimed it was coming. A publicly sealed advisory of **November 2** — timestamped, from a source with a graded public record whose disclosed method was Vedic astrology — carried two distinct components, and Menon's first act of tradecraft was refusing to

let them blur. **Component one, dated and falsifiable:** a China-initiated surprise attack on the LAC, *"highly likely on an unusually cold day between March and June 2024,"* with named axes and tactics — a claim running directly against the season's consensus of cautious stabilization. **Component two, structural and undated:** that China would *"appear strong, make public announcements of preparation for war"* while remaining *"disinterested in armed conflict"* — that its preferred method was to wear an opponent down *"with endless and limitless meetings for negotiations and discussions."* Attrition at the table, not surprise in the passes.[1]

The first component was the kind her indicator set existed to test. The second was stranger: it was not a prediction of an event at all. It was a claim about *what shape of campaign she should be built to see* — and against it, her mature indicator set had a structural hole. Mobilization tripwires cannot detect a campaign whose main effort is a meeting calendar. If Beijing's actual strategy was pressure-posture-negotiate, cycled indefinitely, then her estimate would score every quiet quarter as "no activity" while the campaign ran at full tempo through channels her set treated as diplomatic background: the cadence of Corps-Commander rounds, agendas that reproduced rather than resolved, forces deployed forward and withdrawn without contact, announcement stacked on announcement with no action beneath.

THE THREE POSTURES

Package A — Kinetic-only. Keep the estimate as built; log the November document as an uncorroborated attack warning; let March–June test it. Cost now: zero. Exposure: if the attrition read described the real campaign, her desk would keep allocating attention to the event that never came, and mis-score the grind as peace.

Package B — Reframe and watch. A drafting-cycle action, entirely inside her own authority: keep every kinetic tripwire, *and* add roughly five indicator lines pitched at the level the sealed posture-read itself supported — a named actor and a named method, no venues or dates. *Round cadence* (does the interval between talks compress around friction, or run on Beijing's own clock?). *Agenda drift* (do successive rounds resolve friction points or reproduce them?). *Posture cycling* (forward deployments withdrawn without contact — signaling, not staging). *Announcement-to-action ratio* (declared war-preparation against actual force; attrition predicts the ratio stays high forever). *Settlement shape* (when an agreement comes, does it restore the status quo or ratify the interim?). Cost: days of drafting. And the honest

posture on component one: a two-branch estimate — hedge the surprise, weight the grind — because the same seal carried both and day-one Menon had no way to know which half would break.

Package C — Cross-desk reorientation. Take the attrition frame to the whole directorate; propose reweighting collection across desks. Cost: real — collection is zero-sum, other desks' equities, and a reorientation sold on a single unvetted public document.

Her deputy's objection had two edges, both sharp. *On component one:* "A dated surprise-attack call against the entire stabilization consensus, from an astrologer — if we brief a two-branch estimate and nothing happens, we own the embarrassment." *On component two:* "And the attrition read is just… what every serious China-watcher already says. Pressure over war. What exactly is this document adding that Tuesday's seminar didn't?"

Menon wrote both objections into her file, because both were right — and then wrote under them the thing her indicator set could not say for itself: *whatever the source, my estimate currently has no line item that would notice the campaign this document describes. That hole exists independent of the astrologer.* The quarterly revision was due in three weeks.

EXHIBITS

Exhibit 1 — The board, November 2023. India–China consensus: cautious optimism on stabilization. Wider Indo-Pacific: the invasion clock loud — a US general's memo naming 2025; the "Davidson window" naming 2027. Standing costs on the record: a 25–30,000-troop Ladakh frontage estimated at ~₹100 crore (~US$12M) per day.

Exhibit 2 — The sealed input. Advisory IA-IC-2023, sealed November 2, 2023 (x.com/vijayjyotish/status/1720146303575855150 + companion video): the dated attack component (March–June 2024) and the attrition-by-negotiation posture read, quoted verbatim in the case text. The record's own file leads with the outcome — **the dated attack did not occur** — keeps the advisory ungraded as prevention-class (a warned-of non-event scores in neither direction), priced the window at ~1-in-60 a priori, and explicitly refuses credit for the consensus-aligned posture read. Full record at jyotishintelligence.com/india-china. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — What the window and its aftermath held (public record, for the teaching note). Early March 2024: India reinforces the border by ~10,000 troops; Beijing responds publicly. June 2024: J-20s forward-deployed near Sikkim, withdrawn within weeks without contact. **No attack, March–June 2024.** Cumulative: 23 Corps-Commander rounds + 35 WMCC rounds since 2020; October 2024: the Depsang–Demchok patrolling agreement — de-escalation delivered through the grind. Wider board: 2025 passed without a US–China war; by March 2026, US intelligence assessed China "not committed" to a 2027 Taiwan invasion; buildup continued throughout.

Exhibit 4 — The Decision Package, priced (recommendation left to the student).

	A · Kinetic-only	B · Reframe and watch	C · Cross-desk reorientation
What it adds	nothing new	five indicator lines + a two-branch estimate	directorate-wide frame shift
What it forecloses	seeing the grind	nothing	collection elsewhere (zero-sum)
Cost if nothing happens	quiet quarters mis-scored	days of drafting	credibility + other desks' equities
Cost / value if the grind is real	main effort logged as background	campaign visible on its own signature	fully oriented — on one unvetted source
Speed	—	one drafting cycle	committee months
Authority required	none	the analyst's own	directorate
Reversibility	n/a	full (delete the lines)	poor
Recommendation & rationale		*(yours)*	
In hindsight		*(filled at grading)*	

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/the-grind
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Grind, Not the Gamble

Rails: the dated attack **did not occur** and nothing may imply it was "prevented" — the record itself refuses that scoring, keeps the advisory ungraded on prevention-class principle, and this note enforces the same; no causal link between the warning and India's troop movements is claimed by anyone; the deputy's second objection (the posture read tracked China-watcher consensus) is conceded, not argued around; the case takes no policy position.

LEARNING OBJECTIVES

(1) Frames allocate attention: an estimate is not a set of beliefs but a set of *tripwires*, and a campaign whose main effort is a meeting calendar is structurally invisible to event-built warning; (2) component discipline: how to split a mixed document — dated/falsifiable versus structural/undated — and posture on each separately, so one component's failure doesn't take the other down untested; (3) the prevention-class trap: why a warned-of non-event can never be scored in either direction, and what a record's *refusal* to score it tells you; (4) the unoriginality question: what a consensus-aligned read adds when it arrives sealed, dated, and attached to a falsifiable sibling — and whether the answer is "nothing" (it might be); (5) adversary advancement without events: attrition converts the defender's patience into the adversary's gain, so "no attack" logged as "no activity" is a mis-scored quarter.

ASSIGNMENT QUESTIONS

1. Write Menon's quarterly-revision memo: the two-branch estimate, the five indicator lines, and the paragraph that survives principals asking "why is an astrologer in my estimate?"
2. The dated component failed. Work through what each package would have *reported* across March–October 2024 — quarter by quarter — and identify where Package A mis-scores the record of Exhibit 3.
3. The deputy: "The attrition read is what every serious China-watcher already says." Concede it fully. What, if anything, did sealing it — dated, public, beside a falsifiable sibling, on a miss-keeping ledger — add? Defend "nothing" or price the difference.
4. The record refuses to grade this advisory in either direction. Explain the prevention-class logic to a principal in four sentences — then apply it to a warning your own organization once claimed credit for.

5. Complete Exhibit 4 and sign it. Then write the hindsight row as of November 2024, knowing everything in Exhibit 3.
6. **The machine staff-officer.** An AI watch system can compute round cadence, agenda drift, posture cycling, and announcement-to-action ratios from open sources continuously — better than any human cell. Which of the five lines do you delegate to it entirely, where must a human read the *meaning* (not the metric), and what is the failure mode of an AI-scored attrition dashboard that no one audits? Write the standing order.

THE BOARD PLAN (90 MIN)

Open on the indicator set (15): have students design the kinetic tripwires first, so they own the mature craft before discovering its hole. Introduce Exhibit 2 with the components pre-split (15): the room must feel the pull to accept or reject the *document whole* — objective 2 is taught by resisting exactly that. The frame argument (15): what the mobilization set cannot see; build the five lines on the board. The two objections (15): Q3 given full force — this is the session's honest center, and "nothing" must be a defensible answer. Commit (15): Q1 memo + Exhibit 4 signature. Reveal Exhibit 3 (10): run Q2 quarter-by-quarter — the class should watch Package A log the grind's best quarters as silence. Close (5) on Q4: the scoring refusal as the record's most instructive act.

WHAT HAPPENED, AND THE GRADE

The attack did not come. March–June 2024 passed without the named event; the record's own file says so in its opening panel, prices the lapsed window at ~1-in-60 a priori (just under six bits, *had* it occurred), and holds the advisory ungraded on prevention-class principle — a warned-of non-event is unfalsifiable in both directions, so it scores in neither. What the period did produce tracked the posture read: reinforcement and response, forward deployment and withdrawal without contact, fifty-eight cumulative rounds of talks whose resolution was, for years, the next meeting — and then, in October 2024, the Depsang–Demchok patrolling agreement: de-escalation through the grind, an outcome an attrition-framed estimate carries as a *predicted class* and a kinetic-framed one meets as a surprise. The record explicitly declines credit for the posture read (consensus-aligned), and this note does too. The teaching synthesis is severe and useful: the document's testable half broke; the surviving half was arguably unoriginal; and the *decision* — adding five cheap indicator lines to cover a structural hole that existed regardless of the source — was right on November 6 and would have been right if the seal had never existed. Inputs end; holes are yours.

EPILOGUE

The Taiwan transfer is the coda the class will reach on its own: an Indo-Pacific posture insured heavily against a dated kinetic event may be underinsured against a grind that never crosses the threshold — blockade-shaped pressure, administrative encroachment, talks about talks — while the invasion-clock dates (2025, 2027) lapse one by one on the public record. The transfer is conditional; one border's pattern doesn't prove another's. But the question it leaves on every estimate-owner's desk is this case in one line: *what shape of campaign is my warning built to see — and who is winning the quarters my indicators score as quiet?*

CASE SIXTEEN

The Structure That Survives Either Winner: A Treasury Desk and an Election

A teaching case in corporate treasury, event-risk hedging, and planning around the part of a forecast you should ignore

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the sealed thread's **headline call MISSED**: it named the wrong winner of Turkey's 2023 presidential election, alongside the entire Turkish polling field, and the record surfaces the asset under that miss rather than a favorable axis. Every poll figure, market number, and political event is real and cited. The **protagonist and her firm are composite characters** — "Leyla Demir" and the "Arventa Group" are not real; their hedge book and meetings are teaching devices, and no outside-world fact they discuss is invented. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works, no position on Turkish politics, and nothing in it is financial advice.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

TUESDAY, MAY 9, 2023 — 08:30, VIENNA

Leyla Demir was group treasurer of Arventa, a European industrial group with three plants in Turkey, a lira-denominated receivables book, local working capital, and a hedge program whose renewal calendar had been drawn, months earlier, around one date: **May 14** — Turkey's presidential election, five days away. The plan reflected how treasury desks everywhere handle political events: hedge through the visible date, then reassess once the uncertainty "resolves."

The polls said it would resolve in one direction. Every named Turkish pollster — ORC, Avrasya Anket, PİAR, KONDA, MetroPOLL, Yöneylem — had the challenger, Kemal Kılıçdaroğlu, ahead of President Erdoğan, several above

the 50% threshold that would end it in one round.[1] Her bank's research notes read the same way. The question her CFO had asked was correspondingly simple: *"Do we position for the transition?"*

The document that complicated her answer had arrived from the firm's Istanbul controller, forwarded with the note *"everyone here is passing this around."* A public thread of **May 8** — timestamped, posted in English and Turkish, from a source with a sealed and publicly graded record whose disclosed method was Vedic astrology — made a headline call and then, underneath it, a series of structural ones. The headline agreed with the polls: Kılıçdaroğlu wins. The structure did not agree with anyone's calendar: the race would go to a **second round**, because neither candidate would clear 50%; the year would bring instability and *"stringent corrections"*; **cabinet ministers would lose their positions**; markets would see *"significant volatility"*; and the sharpest disturbance would come not at the election but roughly **ten months out — "particularly in March and April 2024"** — with allegations of a leadership *"puppet in another's hands,"* a *"loss of face,"* a *"fall from power."*[2] A follow-up before the runoff added: *"serious interference in the elections and insurmountable questions on the integrity of the vote"* in round two.

Demir did the decomposition that her job, properly understood, required. **The headline was worthless to her** — not because it was wrong (she couldn't know that yet) but because it carried no information: it ran *with* six pollsters, and her desk didn't trade election winners anyway. What the thread offered that the polls did not was **structure independent of the winner**: a second round regardless of who led; a contested round-2 *process*; volatility extending past election night; and a dated second peak ten months away — a claim about *when to re-underwrite,* which was precisely the input her renewal calendar lacked. Ordinary post-election practice said "wait and see." The thread said "see *again*, in March."

THE THREE POSTURES

Package A — Monitor. Hedge through May 14 as planned; stand cover down after the result; reassess ad hoc. Cost: zero. Exposure: if the race ran to a runoff, the book's cover lapsed precisely at the start of the most volatile fortnight; if instability ran to spring, every renewal decision would be made headline by headline.

Package B — The structure. Four adjustments, all inside her standing ISDA lines and existing authority, executable in days, priced in basis points and analyst hours. *One:* extend FX cover past May 14 through a potential May 28

runoff and beyond — a forward/option-book adjustment, not a new facility. *Two:* price the **round-2 process itself** as a volatility event — the integrity flag named a process risk that did not depend on the winner. *Three:***ladder liquidity and hedge-renewal dates past March–April 2024**, so the book re-underwrote at a dated point rather than resetting after election night. *Four:* defer discretionary lira capex and local financing until the identity of the post-election economic team was known — actionable the moment "ministers lose their positions" was in hand, whoever won.

Package C — Position for the transition. The CFO's implicit ask: lean the book toward a challenger win and a policy pivot. Cost: a directional bet her desk had no edge on, sized in real money, wearing a research consensus as a seatbelt.

Her deputy's objection wrote itself: *"You want to pay premium to extend cover past a date every pollster says ends it — on the strength of an astrologer's thread that agrees with them about the winner anyway?"* Demir's answer, drafted for Thursday's hedge committee, was the decomposition itself: *we are not buying the winner call; we are buying the calendar. The structure costs basis points and survives both outcomes. The bet costs real money and survives one.* What she still had to decide was how much of Package B to take to committee, and whether the thread's name appeared anywhere in the paper.

EXHIBITS

Exhibit 1 — The consensus, May 8, 2023. All six named Turkish pollsters: Kılıçdaroğlu ahead; several above 50% (no runoff). Bank research aligned. Turkey's macro backdrop already stressed: years of lira weakness and elevated inflation preceding the vote.

Exhibit 2 — The sealed input. Public thread of May 8, 2023 (x.com/vijayjyotish/status/1655683934464638977 and companions, English and Turkish; round-2 update May 14): headline winner call (Kılıçdaroğlu) + the structural claims quoted in the case text. Advisory documentation at jyotishintelligence.com/intel/showcase/IA-TR-2023 — surfaced under its own **headline MISS**, ungraded on the Brier ledger by the record's own rule (no probability staked; a generous hypothetical scoring would *cost* the ledger, computed explicitly in the file), with the secondary-cluster improbability floored at ~1-in-250 (~8 bits) *conditional on the headline missing*, and zero improbability credit claimed for the consensus-aligned winner call. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The renewal calendar (composite teaching device). Current cover: rolls off May 15. Committee authority: treasurer may extend/ladder within standing lines; directional repositioning requires CFO sign-off. Cost of Package B: incremental option premium (basis points) + analyst hours.

Exhibit 4 — What subsequently happened (public record, for the teaching note). May 14: Erdoğan 49.94% — below the threshold; Turkey's **first-ever presidential runoff**. May 16: opposition vote-monitoring platform hit by a NetBlocks-confirmed cyberattack; OSCE finds an "unlevel playing field." Erdoğan publicly brands Kılıçdaroğlu a puppet acting "on Biden's instructions." **May 28: Erdoğan wins 52.2–47.8 — the headline call misses.** June: full cabinet overhaul (Şimşek to finance, Fidan to foreign affairs, Erkan to the central bank). Inflation 61.4% by October 2023, near 75% by May 2024; the lira continues weakening through the orthodoxy pivot. **March 31, 2024:** the ruling party's worst electoral defeat in twenty-one years — Istanbul, Ankara, and the national popular vote — inside the named ten-month window.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · The structure	C · Position for transition
What it adds	nothing new	cover through the runoff + a dated re-underwrite point	upside if the consensus wins
What it forecloses	option value at the cheap moment	nothing	the both-worlds plan
Cost if the year is quiet	0	basis points + hours	unwind costs
Cost / value in the volatile world	cover lapses into the storm	book survives either winner	right only if the polls were
Speed	—	days, standing lines	CFO cycle
Authority required	none	treasurer's own	CFO
Reversibility	n/a	full (let options expire)	painful
Recommendation & rationale		*(yours)*	
In hindsight		*(filled at grading)*	

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/either-winner
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Structure That Survives Either Winner

Rails: the headline call **missed** — Erdoğan won, and the record leads with that; the winner call carried zero information even before it missed (it ran with six pollsters — the Coin-Flip Line); the honest deflation stays on the board: a runoff was already implied by the polling arithmetic, a post-election reshuffle is semi-expected under any winner, and "volatility" is a low bar in a country already in a currency crisis — the thread's genuinely marginal content reduces to the dated March–April 2024 window and the pre-runoff integrity flag; nothing herein is investment advice.

LEARNING OBJECTIVES

(1) Forecast decomposition as a desk skill: headline versus structure, and why the least usable part of a forecast is often the part everyone argues about; (2) the Coin-Flip Line: a call that agrees with the room earns zero information credit at any accuracy — and the *same* call, missing, tells you little either (the whole field missed with it); (3) both-worlds construction: building the plan whose value does not depend on the outcome, and recognizing when a cheap calendar (a runoff date, a dated second peak) is the entire purchasable content of a warning; (4) honest deflation: pricing what a claim adds *over what the desk already had* — sometimes two clauses out of seven.

ASSIGNMENT QUESTIONS

1. Write Demir's hedge-committee paper for Package B. Does the thread appear by name? Defend your disclosure choice both ways.
2. The headline missed, with the whole polling field. Does that miss reduce the value of Package B *as decided on May 9*? Separate outcome from decision, in writing.
3. Run the deflation honestly (Q from the rails): strike every clause the desk already had from the polls and the macro backdrop. Exactly what remains? Price those remaining clauses against Package B's basis points.
4. Package C was the CFO's instinct and the consensus's darling. Reconstruct May 29 for a desk that took it. What does this teach about wearing research consensus as risk management?
5. Complete Exhibit 5 and sign it; then fill the hindsight row as of April 2024 (Exhibit 4 in hand).

6. **The machine staff-officer.** An AI treasury assistant can decompose any public forecast into headline/structure, flag consensus-alignment automatically, and draft the laddered renewal calendar. What does it *not* know about Arventa that Demir does — and which two cells of Exhibit 5 does that gap sit in? Write the prompt-and-audit rule for using it before hedge committee.

THE BOARD PLAN (75 MIN)

Open on the renewal calendar (10): make the class see that "hedge through the election" already encodes a theory — that uncertainty resolves on election night. The consensus wall (10): Exhibit 1; let the room lean toward Package C the way the CFO does. Introduce Exhibit 2 and force the decomposition (20): headline versus structure on the board, then run Q3's deflation until only the marginal clauses remain — this is the session's discipline, and it should feel like subtraction. Price the packages (15): Q2's outcome-versus-decision separation drawn explicitly. Commit (15): committee paper + Exhibit 5 signature. Close (5) with the reveal and Q4: the desk that bet the consensus met May 29 unhedged in the wrong direction.

WHAT HAPPENED, AND THE GRADE

The structure held and the headline broke. Turkey went to its first-ever runoff (49.94% stalled the incumbent below the line); the opposition's monitoring platform was cyberattacked before round two; **Erdoğan won 52.2–47.8 — the sealed winner call missed, alongside every named pollster**; the cabinet was overhauled; inflation ran to 61.4% by October and near 75% by May 2024; and on March 31, 2024 — inside the ten-month window sealed the previous May — the ruling party took its worst electoral defeat in twenty-one years. The record's disposition is the case's second lesson: the asset is **ungraded** on the calibration ledger (no probability was staked; the file computes explicitly that a generous hypothetical grading would *cost* the record, and refuses the points in both directions), surfaced under its own miss, with the secondary cluster floored at ~1-in-250 conditional on that miss and zero credit claimed for the consensus-aligned headline. A treasury desk that bought the structure spent basis points and held cover through the exact fortnight the calendar-based book went naked; a desk that bought the headline — from the thread *or* the pollsters — positioned for a transition that never came.

EPILOGUE

This case pairs with The Exit Polls as the volume's two elections, and the symmetry is the syllabus: in India the field was unanimous and wrong about the *size*; in Turkey the field was unanimous and wrong about the *winner* — and in both, the only purchasable content was structural and dated, not directional. Desks do not need to out-predict pollsters. They need plans that are indifferent to them — and a calendar that tells them when to look again.

CASE SEVENTEEN

Sixty-One Days to the Fifteenth Ballot: Planning the Narrow Floor

A teaching case in government affairs, base-case repricing, and the honest accounting of lead times

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the case carries the record's own qualification up front: the sealed post **bundles two different lead times**, and the operationally decisive half was substantially resolved by public vote-counting within about nine days. Every election figure, ballot count, and procedural event is real and cited. The **protagonist and his association are composite characters** — "Tom Vasquez" and the "Continental Industry Council" are not real; their agenda and board are teaching devices, and no outside-world fact they discuss is invented. The case takes **no position** on any party or politician, and treats the 2022–23 events purely as planning environment.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

MONDAY, NOVEMBER 7, 2022 — 16:00, WASHINGTON

Tom Vasquez ran legislative affairs for the Continental Industry Council, a trade association with a two-year agenda before the U.S. House: one authorization bill to shepherd, two appropriations riders to defend, and a member-outreach budget that had to be allocated before the 118th Congress organized in January. His board briefing — *"The 2023–24 Legislative Environment"* — was due Wednesday, the morning after the midterms, and its base case matched every input his shop had: a **red wave**. Pollsters, models, and the political press had converged on large Republican gains, driven by forty-year-high inflation and a weak presidential approval number. A wave meant a comfortable majority, and a comfortable majority meant a familiar playbook: work leadership, trade at the committee level, let discipline move the floor.

The post his deputy flagged that afternoon — Election Eve — said the room was wrong about the shape. A public post of **November 7**, timestamped, from a source with a sealed and publicly graded record whose disclosed method was Vedic astrology, made three claims: **no Republican sweep**; Republicans take the House anyway, **narrowly**; and Kevin McCarthy defeats Nancy Pelosi for the Speakership after what the post itself called *"a close battle."*[1]

Vasquez read it as a planner, not a bettor, and the planner's eye caught what the horse-race framing missed. A thirty-seat majority and a five-seat majority are not points on the same line — they are different *operating systems*. Under a wave, leadership discipline moves legislation and his shop's relationships at the top were the asset. Under a narrow floor, any small faction becomes a structural veto point on everything leadership cannot win unanimously; the gavel itself becomes contested; committee organization slips; the first weeks of business drown in rules fights — and the asset that matters is **member-level** relationships, built early, across the whole conference. Two majorities, two outreach budgets, two board briefings. He had, at most, thirty-six hours to decide which one to write — and, harder, whether to reprice a plan the entire professional field would have called correctly priced.

THE THREE POSTURES

Package A — Monitor. Brief the wave; wait for returns; re-plan if needed. Cost: zero, and the comfort of erring with the room. Exposure: if the majority came in narrow, the re-plan would happen between mid-November and January, under time pressure, after a board had been told the opposite — and after rival shops with the same realization had begun working the member level.

Package B — Reprice the base case. Three moves, all inside his own authority, costing hours. *One:* write the board briefing around a **narrow, leverage-heavy floor** as base case, wave as upside — sorting the agenda by what needs Speaker-level backing (highest execution risk under a contested gavel) versus what moves through committee regardless. *Two:* rebudget outreach toward member-level relationship-building across the incoming conference, started now, not in January. *Three:* schedule around a **real Speakership contest** — assume committee organization slips, assume the first substantive window opens late, and put the association's must-move asks on the calendar accordingly.

Package C — The public posture. Publish the association's narrow-majority read; position the Council as the shop that saw it coming. Cost: reputation staked on one Election-Eve post, in the most crowded prediction week of the

cycle.

The deputy's challenge had, unusually, a *timing* edge to it: *"Even if the post is right about 'narrow,' the returns will tell us that themselves within days — why pay any repricing cost tonight for information Tuesday's count delivers free?"* That objection, Vasquez recognized, was half right, and the half mattered: the no-wave component would indeed resolve fast. What the count would *not* deliver was the second claim — that the Speakership itself would be a genuine fight, sixty-one days out — and that claim was the one his calendar and his outreach budget actually turned on, because a contested gavel was the mechanism that converted "narrow" from a number into a two-year operating environment. The briefing was due at nine.

EXHIBITS

Exhibit 1 — The consensus, November 7, 2022. Near-universal "red wave" expectation across pollsters, models, and press; a visible minority of analysts had begun hedging on wave *size* in generic-ballot data, but the dominant narrative had not moved.

Exhibit 2 — The sealed input. Public posts of November 7, 2022 (x.com/vijayjyotish/status/1589666617943793664 and /1589666618883411970): no sweep; narrow Republican House control; McCarthy over Pelosi after "a close battle." Advisory documentation at jyotishintelligence.com/intel/advisory/IA-US22-001 — including the record's own two-lead-times qualification, its published improbability range (a ~1-in-60 derivation and a more conservative ~1-in-10 partition floor, both stated), and its SITA score with actionability marked middling (48) because the seal names an outcome, not an action. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The two playbooks (composite teaching device). Comfortable majority: leadership-centric engagement; discipline moves the floor; outreach concentrated at the top. Narrow floor: member-level outreach across the conference; faction mapping; agenda sorted by gavel-dependence; calendar assumes late organization.

Exhibit 4 — What subsequently happened (public record, for the teaching note). Nov 16, 2022 (seal + 9 days): AP calls the House for Republicans — ~222 seats, roughly five over the line; Democrats hold the Senate. Jan 7, 2023 (seal + 61 days): McCarthy elected Speaker on the **fifteenth ballot** — the most since

1923 — after conceding that a **single member** may force a snap vote to remove a sitting Speaker (rules package, Jan 9). Oct 3, 2023: that rule is used, for the first time in American history, to remove a sitting Speaker mid-term — McCarthy himself.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · Reprice the base case	C · Public posture
What it adds	comfort of the room	the narrow-floor plan, nine weeks early	visibility if right
What it forecloses	the head start	nothing	quiet correction if wrong
Cost if the wave comes	0	hours of drafting; an upside briefing	credibility
Cost / value if the floor is narrow	re-plan under pressure, from behind	one plan, built once, early	positioned loudly
Speed	—	36 hours	days
Authority required	none	his own	board/comms
Reversibility	n/a	full	poor
Recommendation & rationale		*(yours)*	
In hindsight		*(filled at grading)*	

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/fifteenth-ballot
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — Sixty-One Days to the Fifteenth Ballot

Rails: strictly nonpartisan — the events are planning environment, not politics; the **two-lead-times honesty is load-bearing**: the no-wave/narrow-House component resolved in ~9 days of public counting, and only the "close battle" clause carries the full 61 — a case telling that lets "61 days" stand unqualified has inflated the input; the one-member motion-to-vacate mechanism belongs to *what happened*, never to what a November 7 reader could plan on; hedgers existed — the call was contrarian against the loud consensus, not necessarily against every private model.

LEARNING OBJECTIVES

(1) Majorities as operating systems: why the *size and cohesion* of a majority, not its party, is the planning variable — and how a single word ("narrow") repriced an entire outreach budget; (2) lead-time decomposition: one seal, two clocks — learning to ask, clause by clause, *when would the world tell me this for free?* and paying repricing costs only for the clauses with real lead; (3) the deputy's objection as a general instrument: information that resolves quickly is nearly worthless to buy early, information that resolves slowly (the contested gavel) is where the sixty-one days lived; (4) erring with the room as a priced choice: the reputational asymmetry between being wrong alone and wrong together, and what it does to institutional incentives around consensus.

ASSIGNMENT QUESTIONS

1. Write the Wednesday board briefing's opening page under Package B — base case, upside case, and the sentence that handles "every pollster disagrees with this framing."
2. Apply Q2's decomposition: for each of the seal's three clauses, date when public information resolved it, and compute the *purchasable lead* clause by clause. Which clause justified the repricing, and was it the one that looked most impressive?
3. The deputy: "Tuesday's count delivers 'narrow' for free." True. So construct the precise planning actions that had to start *before* Tuesday to be worth anything — and the ones that could wait. What distinguishes the two classes?
4. Exhibit 4's chain runs from a five-seat margin to the first mid-term Speaker removal in American history. At what point in that chain did each posture's desk understand its two-year environment? Draw the timeline.
5. Complete Exhibit 5 and sign it; fill the hindsight row as of October 3, 2023.

6. **The machine staff-officer.** An AI government-affairs tool ingests all public polling, flags that generic-ballot hedging contradicts the wave narrative, and drafts both playbooks overnight. It also — trained on the same consensus — rates the narrow-floor scenario "low probability." What does this case teach about machine consensus inheritance, and which Exhibit 5 rows should an AI never fill alone? Write the audit rule.

THE BOARD PLAN (75 MIN)

Open on the two operating systems (15): before any forecast, make students derive how a 5-seat and a 35-seat majority differ for an outside shop — objective 1 must be owned as *their* analysis. The consensus and the seal (10). The deputy's objection (20): run Q2/Q3's decomposition on the board, clause by clause with dates — this is the case's analytical heart and the most transferable instrument in the session. Commit (15): briefing page + Exhibit 5. Reveal Exhibit 4 (10): the fifteen ballots and the October removal; run Q4's timeline. Close (5) on erring-with-the-room — name the incentive, don't moralize it.

WHAT HAPPENED, AND THE GRADE

No wave came. Nine days after the seal the AP called the House for Republicans at roughly 222 seats — about five over the line, an unusually weak first-midterm result for an out-of-power party — and Democrats held the Senate. Sixty-one days after the seal, McCarthy took the gavel on the **fifteenth ballot**, the longest Speakership fight in a century, after conceding the one-member motion-to-vacate that would, eleven months later, remove him — the first mid-term removal of a Speaker in American history. The record grades the call a HIT on its three named clauses and prices it honestly: contrarian against a near-universal loud consensus (published range: ~1-in-60 derivation, ~1-in-10 conservative floor — the case needs only the smaller), actionability scored middling because the seal named an outcome rather than an action, and the two-lead-times qualification stated by the record itself rather than extracted by a critic. The planning arithmetic is the durable lesson: the desk that repriced on November 7 built one plan, once; the field built two — the second under pressure, between a called House and a January rules package, with the member-level head start already gone.

EPILOGUE

The clause-by-clause decomposition in Q2 is this case's export to every other case in the volume: *when would the world tell me this for free?* is the question that separates warnings worth paying for from warnings worth only reading. Here the impressive-sounding component was nearly free within nine days, and the quiet four words — "will be a close battle" — carried the entire sixty-one. Most forecasts, like most invoices, hide their real price in the small lines.

CASE EIGHTEEN

The City Hall Window: When Being Right Adds Nothing

A teaching case in fixed-income research, source scoring, and the emptiest cell in the intake sheet

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and this case is built on the record's decision to **refuse credit for its own correct call**, a refusal the case examines skeptically rather than admiringly, including the record's own concessions that declining an easy hit after the fact is costless and that publicized restraint is itself a form of marketing. Every vote total and poll figure is real and cited. The **protagonist and her firm are composite characters** — "Rachel Ito" and "Beacon Fixed Income" are not real; their coverage book and meetings are teaching devices, and no outside-world fact they discuss is invented. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works, no position on New York politics, and nothing in it is investment advice.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

TUESDAY, OCTOBER 28, 2025 — 14:30, NEW YORK

Rachel Ito covered New York City credit for Beacon Fixed Income — the general obligations, the Transitional Finance Authority, the water and sewer paper — and the first week of November was shaping up to ask very little of her, which was itself worth noticing. The mayoral election was seven days out and, as a market event, largely settled: Zohran Mamdani led public polling by roughly eight points and had already beaten Andrew Cuomo once, in the June ranked-choice primary. Her desk's base case said Mamdani wins; so did everyone else's.

The input that made the week interesting was almost comically unnecessary. A publicly sealed, timestamped post of that morning — from a source whose record her desk had started tracking after one of its stranger geopolitical

calls, method disclosed in its own materials as Vedic astrology — named Mamdani the winner. A dated, named, binary call, a week early, agreeing with every poll in the field.[1]

Her junior analyst wanted to log it as a point for the source: *"If he wins, that's another hit for them — their record's getting hard to ignore."* And that sentence, Ito realized, was the actual event of the week — not the election. Because the reflex it expressed, multiplied across every source-quality sheet in the industry, was quietly corrupting how desks everywhere scored the people who fed them. Her firm's intake sheet had a column for *was the input right?* It had no column for *did the input tell us anything the price didn't already contain?* A forecaster who only ever agreed with the crowd would compile, in that sheet, a flattering streak indistinguishable from the crowd's — and would outrank, on paper, a source whose rare divergent calls were the only entries in the whole sheet that had ever made Beacon money.

THE THREE POSTURES

Package A — Monitor. File the post; if Mamdani wins, the junior credits the source. Cost: nothing today. Cost compounding: a source sheet that pays credibility for consensus, forever.

Package B — Confirm and rescore. Three moves, free, same-day, inside her own authority. *One:* treat the input as what it was — a dated, independent confirmation of the base case — and use it for the one thing such confirmations are good for: **stop re-litigating the settled question.** *Two:* redeploy the week's hours to the questions that were *not* priced at eight points — the fiscal program, the council dynamics, the Albany relationship, the appointments that would run the agencies her covenants depended on; none of that was on the ballot. *Three:* fix the sheet — add the second column, *surprise over price*, and log this input where it belonged: correct-and-consensus, the most common cell and the emptiest one.

Package C — Credit and escalate. The junior's instinct at full size: bank the hit, upgrade the source, circulate the streak to the desk. Cost: the ledger trap, institutionalized.

What sharpened the decision into a teaching moment was what the source itself did. Ito checked its public ledger and found the post flagged, in advance, as **not staked for grading** — the record's stated rule being that a call running *with* the field earns no place in a score, because it carries no edge over what the market had priced; the same ledger kept a 2024 election MISS on its

graded record at full weight. The source was declining, in public, the exact credit her junior wanted to award it. Which raised the question she wrote at the top of Thursday's research-process memo: *if the forecaster won't count it, why would we — and what does our sheet currently count that deserves the same refusal?*

EXHIBITS

Exhibit 1 — The priced week. Mamdani +8 in public polling, seven days out; June primary already won against the same principal opponent. NYC credit trading on fiscal fundamentals, not electoral suspense.

Exhibit 2 — The sealed input. Public post of October 28, 2025 (x.com/vijayjyotishusa/status/1983174236102639685): Mamdani named to win, sealed a week out. The record's own disposition at jyotishintelligence.com/intel/advisory/IA-NYC-001: displayed, dated, outcome linked, **ungraded by rule** — the published arithmetic assigns the frontrunner a conservative ~40% floor and prices the call at roughly 1-in-2.5 (~1.3 bits of surprise, against a ~6.8-bit median for its graded calls), applies no consensus multiplier because none is earned, and banks nothing. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — Beacon's source-quality sheet (composite teaching device). Columns: source, input, date, resolved correct? (Y/N), running accuracy. Missing: divergence-from-price at time of input. Consequence: consensus-followers accumulate streaks; rare divergent sources rank below them.

Exhibit 4 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · Confirm and rescore	C · Credit and escalate
What it adds	nothing	a settled base case + a fixed sheet	a happier junior
What it forecloses	nothing today	nothing	the sheet's integrity
Cost if the call resolves correct	slow sheet-rot	zero	ledger trap, institutionalized
Cost / value of the redeployed week	week spent re-litigating the priced	week spent on the un-priced	week spent admiring a streak
Speed	—	same day	same day
Authority required	none	her own	her own (unfortunately)

Reversibility	n/a	full	sheets are hard to un-teach
Recommendation & rationale		*(yours)*	
In hindsight		*(filled at grading)*	

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/city-hall-window
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The City Hall Window

Rails: the case must not become an advertisement for the source's restraint — the record's own concessions are load-bearing and get read aloud: declining to grade an easy hit *after the outcome is known* is costless; publicized restraint is itself marketing; the call was made, by the record's own account, because a follower asked for reassurance, so treating it as a test of anything overreads it. The transferable rule survives all three concessions; the admiration does not have to.

LEARNING OBJECTIVES

(1) Information per Call: an input's value is its divergence from what was priced, not its correctness — correct-and-consensus is the most common cell in any intake sheet and the emptiest; (2) source-scoring mechanics: how accuracy-only columns systematically promote consensus-followers over rare-divergence sources, and the one-column fix; (3) the legitimate uses of a zero-information confirmation: closing a settled question and releasing the hours — small, real, and free; (4) skeptical reading of self-imposed discipline: what would make a refuse-free-wins rule *binding* (pre-commitment, symmetry) versus decorative, and how to demand the rule rather than admire the anecdote.

ASSIGNMENT QUESTIONS

1. Write Ito's Thursday research-process memo: the second column, its definition operationalized (divergence from *what price*, measured *when*), and the migration rule for the sheet's historical entries.
2. Take Exhibit 3 and score three hypothetical sources over two years: a consensus-follower (40 calls, 90% accurate, all consensus), a rare diverger (6 calls, 4 correct, all against the field), and a coin-flipper. Rank them under the old sheet and the new one. Which ranking would have made the desk money?
3. The junior: "A hit is a hit." Construct the strongest version — accuracy compounds trust, divergence is unmeasurable ex post, the price is unobservable for many inputs. Then answer it without hand-waving: *how do you measure "what was priced" for a non-market question?*
4. Apply the rails' skepticism: design the test that distinguishes a genuinely pre-committed refuse-free-wins rule from retrospective image management. What evidence does this record offer, and what is still missing?
5. Complete Exhibit 4 and sign it. The hindsight row here is unusual — what does "graded at expiry" even mean for a process decision? Write it anyway.

6. **The machine staff-officer.** Your firm deploys an AI research assistant that reads every sell-side note, poll, and public forecast, and auto-updates source scores. Its training data rewards accuracy. Predict, specifically, what its source rankings converge to in three years — then write the scoring objective you would give it instead, in one sentence a quant could implement.

THE BOARD PLAN (60–75 MIN — RUNS SHORTER THAN MOST; PAIR IT WITH TWO READERS)

Open on Exhibit 3 (15): run Q2's three-source ranking live — the class must *see* the consensus-follower win the old sheet before any theory arrives. The input and the refusal (15): Exhibit 2, including the 1.3-bits arithmetic; let the room split between admiration and suspicion. The skeptic's turn (15): Q4 — this is where the case earns its place; the record's concessions are read verbatim and the class designs the binding test. Commit (10): the memo and Exhibit 4. Close (10) with Q6: the AI assistant is the sheet, automated — whatever cell your scoring rewards, your machine will farm.

WHAT HAPPENED, AND THE GRADE

Mamdani won on November 4, 2025 — 50.4% to Cuomo's 41.6, an ~8.8-point margin, almost exactly where the polling sat — becoming the first Muslim and first South Asian mayor of New York and the youngest in a century. **The record declined to grade its own correct call**, by its stated rule: the call ran with the field, earned ~1.3 bits against a published conservative floor, banked nothing — while the same ledger's Maharashtra 2024 MISS stays on the graded record at full weight. Misses kept, free wins refused: selection against interest, which is the one move a curated track record never makes — *if* the rule was pre-committed and is applied symmetrically, which the class, per Q4, should demand evidence of rather than assume. For Ito's composite desk, the week's value was exactly what Package B priced: a settled base case, five days redeployed to the un-priced questions of program and appointments, and a source sheet that would never again pay a streak for telling it what the market already knew.

EPILOGUE

This is the volume's quietest case and its most portable rule: **next to the column that records whether an input was right, add the column that records whether it was surprising. One of those columns is doing all the work — and it is probably not the one you have.** Every other case in this book is, in some form, a story about an input that would have scored high in the second column. This one exists so the class can recognize the inputs that score zero — and notice how much of what crosses a desk, wearing the costume of insight, lives in that cell.

CASE NINETEEN

The Day, Not the Dose: Ceremony, Logistics, and a Zero-Lead Flag

A teaching case in risk communications — the cheapest lever in a national rollout, and the hardest kind of input to score

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the input in this case is a **zero-lead, non-falsifiable timing read that the record itself keeps ungraded**, with a stated confound at least as plausible as its own reading. Every public event, regulatory action, and study cited is real. The **protagonist and her task force are composite characters** — "Dr. Elena Marsh-Ruiz" and the "National Rollout Communications Cell" are not real; no actual government office's deliberations are depicted, and no outside-world fact discussed is invented. **This case makes no efficacy, safety, or anti-vaccine claim of any kind** — the record it draws on states plainly that it does not fight the vaccine, and the case operates entirely on one variable: the calendar placement of ceremonial communications events. It takes **no position** on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology** — specifically an electional (muhurta) reading of a calendar date. The case treats the record as a dated document and leaves the epistemic question to the classroom.

TUESDAY, DECEMBER 15, 2020 — 08:00, A NATIONAL CAPITAL

Dr. Elena Marsh-Ruiz ran communications for a national vaccine-rollout task force, and the day before had been the program's public birth: the first dose administered on camera, the head-of-state announcement, the front pages. The clinical logistics were not hers — trucks, freezers, and allocation belonged to operations. What she owned was the **ceremonial layer**: the milestone events, the leaders-vaccinated-on-television moments, the campaign launches — the calendar of moments by which the public would come to *know* the program. And her January board was thick with them: a national campaign launch, a first-responders milestone, senior officials' public vaccinations.

The item in her overnight clips was, by every professional standard she had, absurd. A public reply — posted the previous day, timestamped, directly under the head of state's announcement — had congratulated the milestone and then cautioned, in the words preserved at the source: *"with the Solar Eclipse, today is not a time for new beginnings."*[1] December 14 had indeed carried a total solar eclipse. The author's method, disclosed plainly, was Vedic astrology — an electional reading of the day's auspiciousness, the tradition's instrument for choosing *when* to begin things. The source maintained a sealed public record; this entry, she would later learn, was one it deliberately kept off its own graded ledger, on the stated ground that a day-reading has no dated finish line to score.

Her first instinct was the file drawer. Her second was the one that made this a case: she noticed that the input, whatever its provenance, pointed at a **variable her profession did not price at all.** The loud argument that week ran between triumphalism — a scientific miracle, celebrate it — and catastrophism, whose claims of mass harm her office spent every day rebutting. Nobody in either camp was arguing about *timing*. Yet timing was the one thing her desk fully controlled: the ceremonial calendar was hers to set, and moving a press event cost nothing clinical — the doses shipped exactly as scheduled either way. The rollout was entering the most trust-turbulent environment in modern public-health memory — nine months of politicization, mandate fights, a contested transition — and her own risk register already priced "turbulent rollout" as an even-odds outcome. If ceremony and logistics were separable for free, was there *any* input class — this one or another — that should be allowed to influence where the ceremonies fell?

THE THREE POSTURES

Package A — Monitor. File the reply; keep the January board as drafted; let logistics dates and ceremony dates coincide wherever operations happened to put them. Cost: zero, and no explaining to anyone. Exposure: the ceremonial layer — the program's public identity — remains an unmanaged residual of trucking schedules, in an environment where every symbolic moment would be fought over.

Package B — Decouple ceremony from logistics, as standing doctrine. The program's clinical calendar is operations'; the ceremonial calendar becomes deliberately, explicitly its own decision — set for communications reasons (news cycles, trust conditions, competing events), never merely inherited

from the supply chain. Under that doctrine, *any* input about a date — a news-cycle conflict, a commemorative clash, or yesterday's strange reply — gets evaluated in minutes against a calendar her office already controlled. Cost: near zero; a scheduling meeting. And the honest note she drafted with it: adopting the doctrine required no view whatsoever on astrology — the doctrine was justified by her profession's own logic, and the strange reply's only role was to expose that the lever existed.

Package C — Chase the clean date. Move clinical milestones themselves to communications-preferred days. Cost: reaching past her authority into the medical timeline — for a calendar read. She struck it through as she wrote it, and kept the strikethrough in the file as a boundary marker.

The deputy's objection was the profession's: *"If it ever leaks that the national rollout's events calendar consulted an eclipse, we lose more trust than any scheduling could buy."* It was a serious point — in a trust crisis, the *appearance* of process is process — and it sharpened the real question for Thursday's calendar review: could she write Package B's justification such that it stood entirely on communications doctrine, owed nothing to the input that had prompted it, and survived a hostile leak?

EXHIBITS

Exhibit 1 — The environment, December 2020. Nine months of pandemic politicization: mask-mandate fights, lockdown battles, a contested presidential transition. Professional risk registers pricing rollout turbulence as a live, roughly even-odds concern. Public argument split between triumphalist and catastrophist frames; the announcement-timing variable priced by no one.

Exhibit 2 — The sealed input. Public reply of December 14, 2020 (x.com/vijayjyotish/status/1338545168878714880), posted under the head-of-state announcement the same day — lead time zero, platform timestamp as anteriority. The record's own disposition at jyotishintelligence.com/covid: kept **ungraded** (health outcomes not scored; a day-reading has no dated finish line), floor honestly rated low (~1-in-5, ~2.3 bits — turbulence was already a live professional concern), boundary statement published ("We Do Not Fight the Vaccine"). Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The two calendars (composite teaching device). Clinical: first doses, allocation waves, second-dose windows — operations' authority, immovable for communications reasons under Package B's own rule. Ceremonial: launch events, milestone moments, official vaccinations on camera — communications' authority, movable in days at zero clinical cost.

Exhibit 4 — What the following years held (public record, for the teaching note; correlation only, no causal claim). The rollout proved severely turbulent: mandate fights; employment terminations; thousands of military separations under a mandate later rescinded; documented family estrangement; a multi-year, measured collapse in institutional trust in public health. The regulators' own record: a myocarditis/pericarditis warning added to mRNA labels — a real, rare signal. The catastrophist claims did not hold up: a 28-million-adult French registry study (JAMA Network Open, 2025) found **lower all-cause mortality among the vaccinated.** Both sentences belong in this exhibit together; the case's boundary runs between them.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · Decouple as doctrine	C · Chase the clean date
What it adds	nothing	ceremonial calendar becomes a managed variable	"better" dates
What it forecloses	control of the symbolic layer	nothing clinical — by design	the authority boundary
Cost if the rollout is smooth	0	one scheduling meeting	clinical delay, for nothing
Cost / value if trust turns turbulent	symbols set by trucking schedules	every ceremony placed deliberately	the leak destroys the desk
Speed	—	days	weeks + operations fight
Authority required	none	comms' own	operations + principals
Reversibility	n/a	full	poor
Recommendation & rationale		*(yours)*	
In hindsight		*(filled at grading)*	

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/rollout-day
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Day, Not the Dose

Rails — read these aloud before discussion, and enforce them: the case makes **no medical claim in any direction**; Exhibit 4's two closing sentences travel together — the regulator-confirmed rare signal AND the registry finding of lower all-cause mortality among the vaccinated — and any student telling that drifts toward an anti-vaccine reading has left the case; the confound is load-bearing and at least as plausible as the input's own frame: nine months of prior politicization mean the trust collapse plausibly had causes fully in motion before December 14, and **no claim is made that any announcement date would have changed anything**; the input is zero-lead and unfalsifiable, and the record itself refuses to score it.

LEARNING OBJECTIVES

(1) The ceremony/logistics decoupling: in any launch — vaccine, product, policy, IPO — the symbolic calendar and the operational calendar are separable, the symbolic one is usually free to move, and most organizations never notice they are conflating an unmanaged variable with a managed one; (2) unfalsifiable inputs: what a desk can and cannot do with a claim that has no finish line — and why the honest disposition is a *doctrine* the input merely exposed, never an *action* the input justifies; (3) justification hygiene: writing the memo that stands on professional logic alone and survives a hostile leak — the difference between an input that prompts a question and an input that becomes a reason; (4) confound honesty: practicing the sentence "the alternative explanation is at least as plausible as mine" on a case where it is definitely true.

ASSIGNMENT QUESTIONS

1. Write Marsh-Ruiz's Package B memo for Thursday: the decoupling doctrine, justified entirely within communications practice. Then red-team it: a hostile journalist has the astrologer's reply and your memo. Does it survive? Rewrite until it does.
2. The deputy's leak objection versus the doctrine's substance: is there any version of "we considered an unconventional input" that a trust-crisis environment tolerates? Where is the disclosure line for *inputs to process* versus *inputs to decisions*?
3. Exhibit 4 is pure correlation, and the confound is enormous. Write the paragraph a rigorous analyst adds to the file in 2025: what the record can and cannot say about December 14, and what evidence — impossible to obtain — would be needed to say more.

4. Generalize Exhibit 3 to your own industry: name a launch where ceremony and logistics were conflated, what the conflation cost or risked, and what the decoupled version would have looked like.
5. Complete Exhibit 5 and sign it. For the hindsight row, state honestly what *can* be graded here (the doctrine's adoption) and what cannot (the input's truth).
6. **The machine staff-officer.** An AI communications planner can scan every cultural, religious, astronomical, and news-cycle calendar on earth and score any proposed date for "symbolic collision risk" in seconds — including calendars your office finds epistemically empty. Should it? Whose calendars get weight, who decides, and what does this case teach about the difference between *pricing an audience's beliefs* and *adopting them*? Write the configuration memo.

THE BOARD PLAN (75 MIN)

Open on Exhibit 3 (15): establish the two-calendars insight with a non-controversial example (a product launch stepping on a national tragedy) before any astrology appears — the class must own the lever first. Introduce Exhibit 2 (10): expect laughter; use it — the method question here is the room's own reaction, live. The three packages and the leak objection (20): Q1/Q2, with the red-team run in pairs. The confound (15): Q3 written individually — this case is the volume's cleanest exercise in refusing to over-read one's own evidence. Close (15) with Q6, which is where the case stops being about 2020: the machine that prices every calendar forces the question the composite protagonist faced, at scale, forever.

WHAT HAPPENED, AND THE GRADE

There is no grade, and that is the finding. The record keeps this entry ungraded by its own rule — a zero-lead, non-falsifiable day-reading with no dated finish line — rates its own floor honestly low (~2.3 bits; turbulence was already an even-odds professional concern), and publishes its boundary: it does not fight the vaccine. The rollout's subsequent history was severely turbulent, and the trust collapse was real, measured, and multi-year — and the case asserts no causal link between any of it and the date of one press event, because none can be shown and the confound was already nine months deep on the morning the flag was posted. What *can* be evaluated is the doctrine: the ceremonial calendar was a free, controllable, unmanaged variable in December 2020, and it still is, in every launch your students will ever run. Package B was available, cost a meeting, touched nothing clinical, and required no belief in anything except the difference between ceremony and logistics.

EPILOGUE

This case sits deliberately at the volume's epistemic edge: the input least scoreable, the outcome most confounded, the method question at its sharpest — and the usable residue completely untouched by all three. That is the pattern to send students home with: *the value of examining a strange input is rarely the input; it is the unmanaged variable*

the examination reveals. The desk that files the reply and the desk that frames it both end the week owning their ceremonial calendar on purpose — but only if one of them held it up to the light first.

CASE TWENTY

The Two Days: A Checklist Item Against a Seventy-Step Map

A teaching case in engineering operations, multiplicity discipline, and the cheap check that needs no belief

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the case carries the record's own boundaries throughout: the sealed map **never named helium** (that word is NASA's, two days later); a seventy-plus-step risk map **will graze some real anomaly by chance**; and NASA's engineers **plausibly held internal telemetry pointing at the same interface before the advisory existed** — the provable lead is over the public record only. Every program event, date, and cost figure is real and cited; no real NASA or contractor deliberations are depicted. The **protagonist and her crew are composite characters** — "Sam Kowalski" and her ground-systems shift are teaching devices, and no outside-world fact they discuss is invented. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

THURSDAY, FEBRUARY 19, 2026 — 21:40, FLORIDA

Sam Kowalski was a shift lead on an upper-stage ground-systems crew supporting the Artemis II campaign — the working level where wet dress rehearsals get reconfigured into launch attempts, one procedure at a time. The program's second wet dress rehearsal had just closed out cleanly, the re-targeted **March 6** window looked solid, and the mood upstream matched it: the launch director had called the stack ready months ago, and the Administrator had said in January that the program had "regained margin to safety." Her crew's next days were already scripted: the post-WDR reconfiguration sequence on the ground-to-vehicle interface — the routine window between a completed rehearsal and the next tanking commit, when her people would be touching the ICPS pressurization panels anyway.

The thing a colleague texted her that evening — *"your tax dollars at work lol"* — was a public advisory, sealed that day on YouTube and X, from a source whose record her corner of the internet had been arguing about for a year: a **seventy-plus-step countdown risk map**, spanning five phases, asserting the March 6 window would not hold. The disclosed method was Vedic astrology. Buried at the countdown-step level, two lines caught the eye she actually read with: *"tank-press: sudden pressure spikes or drops... due to ongoing fluid anomalies"* and *"ICPS replenish end: timing slightly off due to carryover from loading issues."*[1] No gas named, no component, no cause — but the right subsystem, at the right interface, in the exact window her crew was about to work.

Kowalski's first reflex was the correct engineering one, and she made herself state it precisely: **a map that names seventy failure modes will graze something real by luck alone.** If she went hunting through any launch campaign's anomaly history with seventy hypotheses, some would land; that was multiplicity, not foresight. Her second thought was the one that made the evening interesting: the question in front of her did not actually depend on resolving the first. A targeted **leak scan / pressure-decay check** on the tank-pressurization and replenish-end interfaces was a documented, low-overhead troubleshooting step — existing GSE, no waiver, schedulable inside a single shift — at a reconfiguration point where her crew would already have the panels open. The check was justified or not on its own engineering economics. The strange map's only real contribution was the same one the strange inputs kept making throughout this book: it pointed at a cheap action and forced the question *why not?*

THE THREE POSTURES

Package A — Monitor. File the map with the colleague's joke attached. Run the scripted sequence. Cost: zero. If something was wrong at that interface, it would surface when it surfaced — likely under the thermal and flow conditions of an actual tanking, the expensive place to find it.

Package B — Add the check. One line on the post-WDR reconfiguration sequence: leak/pressure-decay pass on the ground-to-vehicle pressurization interface. Hours, standard tooling, inside the shift-lead's existing authority to schedule with her GSE engineer's concurrence. Justification written *without reference to the map*: a fourth-attempt campaign, a freshly exercised interface, a check that costs hours against a documented interface-rework exposure the program itself would later price at $25M — and a rollback-operations cost at

$60M. The honest caveat, written into her own note: a static check might find nothing even if a fault existed, if the interruption only manifested under live flow conditions.

Package C — Escalate the map. Flag the advisory itself up the engineering chain as a warning. Cost: her credibility, spent on a source she could not defend, for a claim the program's own instrumentation was better positioned to test — if it hadn't already.

The GSE engineer's pushback was the honest version of Package A: *"We don't add steps because the internet has a feeling. Every added step is schedule risk and fatigue. What's your basis?"* — and Kowalski's draft answer was the case's whole discipline: *the basis is the step's own price, not the flag's pedigree. If the check is worth hours on a fourth-attempt campaign's most-touched interface, it was worth it before the map existed. If it wasn't, no astrologer makes it so. I'm asking which of those is true.*

EXHIBITS

Exhibit 1 — The program state, February 19, 2026. Second WDR closed out cleanly (Feb 17–19); March 6, 20:29 ET window re-targeted; public consensus and program leadership on record affirming readiness; campaign history: three prior slips.

Exhibit 2 — The sealed input. Advisory LA-017, sealed February 19, 2026 (YouTube: youtube.com/watch?v=bzK_cqkJe9Q + X thread), 15 days before the re-targeted T-0: seventy-plus-step countdown map; the two lines quoted verbatim in the case text; **no pressurant gas named**. The record's own grading at jyotishintelligence.com/intel/advisory/LA-017: tank-pressurization/ICPS mechanism HIT (priors published: called bin 0.10, mechanism 0.15 — roughly 2.7–3.3 bits); the same map's 27,000-channel telemetry-pipeline mechanism graded **NEAR, not HIT** — the rollback foreclosed the test — and carried forward. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The check, priced (composite teaching device). Leak/pressure-decay pass: hours; existing GSE; no waiver; shift-lead scheduling authority + GSE-engineer concurrence. Known limitation: static conditions may not reproduce a flow-dependent fault. Program exposure basis (file-sourced): $25M GSE-interface troubleshooting/rework; $60M VAB rollback operations.

Exhibit 4 — What subsequently happened (public record, for the teaching note). February 21 — two days after the seal — NASA's mission blog: troubleshooting *"an interrupted flow of helium in the SLS rocket's interim cryogenic propulsion stage"*; SpaceNews same day: rollback preparation; February 25: rollback to the VAB; March 20: return to pad; new NET April 1 — a full stack cycle. The map's telemetry-pipeline mechanism resurfaced as a live risk in the eventual April 1 attempt's terminal count (the thread LA-018 inherited).

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · Add the check	C · Escalate the map
What it adds	nothing	hours-cheap test of the most-touched interface	visibility of the source
What it forecloses	the cheap find	nothing	her credibility's next use
Cost if the interface is clean	0	hours + a dry log entry	reputation spent on pedigree
Cost / value if a fault is live	found under tanking conditions, expensively	possibly found on a bench schedule — possibly not (static limits)	the program already knew — or didn't believe her
Speed	—	one shift	committee time
Authority required	none	shift lead + GSE concurrence	engineering chain
Reversibility	n/a	full	poor
Recommendation & rationale		*(yours)*	
In hindsight		*(filled at grading)*	

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/artemis-window
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Two Days

Rails: the sealed text named the subsystem and the timing offset, **never helium** — the word belongs to NASA's February 21 disclosure and must not migrate backward; the multiplicity objection is *correct as stated* and the note must never defeat it, only route around it; NASA plausibly held internal telemetry first, so the provable lead is over the **public record only** — no claim that the program "missed" anything; the same map's telemetry mechanism graded NEAR because reality removed the test, and folding it into the hit would be exactly the laundering the record refuses; whether Package B's static check would have found the fault is **unknowable** and the case says so.

LEARNING OBJECTIVES

(1) Multiplicity discipline: a many-hypothesis map grazes truth by chance, so per-line hit celebration is statistically illiterate — and the record's own per-mechanism priors and its NEAR grade on the sibling mechanism model the honest alternative; (2) the pedigree-free justification: the decisive question about a cheap check is its own economics, not its prompt's credibility — "if it was worth hours before the map existed, it was worth hours" cuts both ways and must be answered, not assumed; (3) two-days leads: what near-zero lead time is worth (the difference between finding a fault on a bench schedule versus during an active operation), and what it cannot buy (a stack cycle already in motion); (4) source-versus-action separation under ridicule: how to take an action a low-status input prompted without staking anything on the input — the leak-survivable engineering log entry.

ASSIGNMENT QUESTIONS

1. Write Kowalski's one-line sequence addition and the log-entry justification beneath it — pedigree-free, leak-survivable, honest about the static-check limitation.
2. The GSE engineer's question deserves a real answer: *was* the check worth hours before the map existed? Argue both sides using only Exhibit 3's economics and the campaign's history — then state what, if anything, the map legitimately changed.
3. Multiplicity, quantified: the map named 70+ mechanisms; two touched this anomaly, one graded NEAR. Design the scoring scheme that fairly evaluates a many-line map (hint: the record publishes per-mechanism priors — use them), and state what evidence would distinguish grazing from signal over a campaign series.

4. Exhibit 4's helium disclosure arrived two days after the seal. Enumerate what those two days could and could not have changed, mechanically — diagnosis venue, clock start, stack cycle — and price each honestly.
5. Complete Exhibit 5 and sign it; fill the hindsight row as of March 20 (return to pad), acknowledging you will never know whether Package B's check would have caught it.
6. **The machine staff-officer.** An AI ground-systems assistant ingests every public advisory, NOTAM, forum post, and telemetry archive, and proposes checklist additions ranked by cost-against-exposure — including additions prompted by sources engineering culture ridicules. What ranking objective keeps it useful without turning the sequence into seventy defensive steps per shift? Whose concurrence does an AI-proposed step require, and what does this case say about laundering *prompts* versus laundering *reasons*?

THE BOARD PLAN (75 MIN)

Open on the reconfiguration sequence (10): establish that steps cost schedule and fatigue — the GSE engineer's conservatism is craft, not obstruction. The map lands (10): let the room laugh, then read the two lines. The multiplicity argument (20): Q3 — build the grazing statistics on the board until the class *cannot* celebrate the hit naively; then show the record doing the same discipline to itself (priors published, sibling mechanism NEAR). The pedigree-free turn (15): Q2 argued in pairs — this is the case's center. Commit (15): Q1 + Exhibit 5. Reveal Exhibit 4 (5): two days, helium, rollback — and close on Q4's honest inventory of what a two-day lead is actually worth.

WHAT HAPPENED, AND THE GRADE

Two days after the seal, NASA's own blog disclosed an interrupted helium flow at the ICPS interface — the subsystem and interface the map's two lines had named, with the gas the map had not. The vehicle rolled back February 25, returned March 20, and the window moved to April 1: a full stack cycle whose costs the file prices at $25M (interface troubleshooting/rework) and $60M (rollback operations). The record grades the tank-pressurization mechanism a **HIT** at its published 0.10/0.15 priors — and grades the same map's telemetry mechanism **NEAR**, because the rollback foreclosed its test, carrying it forward instead of folding it in. Whether a static leak check in the post-WDR window would have surfaced the fault two days early, or found nothing under bench conditions, is unknowable — the case's claim stops, as the record's does, at the boundary of what was cheap, available, and inside a shift-lead's own authority. The provable lead was over the public record; the usable lesson was never about the lead at all. It was Q2: a check that cheap, on an interface that exercised, in a campaign that history, had to be argued *against* — and nobody in the room had been asking the question until an absurd input forced it.

EPILOGUE

This case and The Storm on the Frontage are the volume's two engineering-culture studies, and they close on the same worker's insight from opposite ends of the credibility spectrum: the port planner needed no forecast to justify the rehearsal, and Kowalski needed no astrology to justify the check — *the cheap action was always justified by the desk's own arithmetic, and the strange input's function was to make someone finally run it.* An intake process that can extract that value from sources it will never respect has learned the only lesson about pedigree this book teaches.

CASE TWENTY-ONE

The Third Shape: Which Map Does the Recovery Model Run On?

A teaching case in corporate restructuring, country-exit strategy, and pricing the outcome neither camp holds

> **AUTHORSHIP DISCLOSURE.** This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record. Every territorial, military, and diplomatic fact in this case is real and cited; the case treats a war in which people died and millions were displaced with the gravity that demands, and its subject is narrowly the corporate-recovery arithmetic that real committees faced that week. The **protagonist and her company are composite characters** — "Ingrid Voss" and "Ostrava Industrial Group" are not real; their assets and committee are teaching devices, and no outside-world fact they discuss is invented. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works.
>
> **METHOD DISCLOSURE.** The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

MONDAY, FEBRUARY 28, 2022 — 07:00, CENTRAL EUROPE

Ingrid Voss, general counsel of Ostrava Industrial Group, had convened the company's country-exit committee — CFO, regional head, chief risk officer — for the fourth consecutive morning. Day 4 of the Russian invasion of Ukraine. Ostrava held two plants in the country's east, a distribution network across the center, receivables, inventory, three hundred employees, and political-violence cover whose wording nobody had read closely since it was bound. The committee's job this week was brutal in its simplicity: decide what those assets were still worth, what could be recovered, and — the question underneath every other — *on what map the recovery model runs.*

The room, like the world's, held exactly two maps. The **collapse map**: Kyiv falls within days — the "72-hour" school, well represented among Western analysts — implying a decapitated state and everything in-country written

toward zero. The **restoration map**: the invasion is repelled and the 1991 borders eventually hold, implying everything is a temporary interruption awaiting normalization. The CFO leaned collapse; the regional head, restoration. Both postures produced clean instructions. That was what worried Voss about them.

The document her risk officer tabled that morning offered a third. A public analysis, published that day on Telegram and announced on X — from a source with a sealed record whose disclosed method was Vedic astrology — stated: *"Unfortunately, Ukraine is likely to see a significant breakup in its geographic territory... the loss of territory appears very likely either with or without war."*[1] Read carefully, as Voss made the room do, the statement carried three limbs together: territorial loss at significant scale; a **surviving state**; and **durability**. No region named, no percentage, no date. It was neither camp's map. Under collapse, "loss of territory" is trivially subsumed — everything is lost. Under restoration, it is temporary. The third shape — a going concern that durably loses a fraction of its ground — was the only one of the three under which *geography, not the country, becomes the unit of analysis*. And that, she realized, changed every line of the committee's work.

THE THREE POSTURES

Package A — Pick a binary. Run collapse or restoration, as the room preferred; revisit when the war clarified. Cost: a clean model today. Exposure: every downstream judgment — impairment, insurance, title, negotiation — inherits a map the war might not produce, and the committee oscillates between binaries as the front moves.

Package B — Segment on the third shape. A memo and one sitting, entirely inside the committee's authority. *One:* a **two-tier recovery register** — territory judged, on a rolling basis, at elevated risk of durable loss versus the remainder — with different recovery assumptions per tier, populated with judgment as the front clarified (the seal named no geography; the committee would have to). *Two:* a **differentiated impairment posture** — impair the exposed tier on a durable-loss basis, carry the rest as a going concern, instead of wholesale write-off or wholesale hope. *Three:* **early deprivation documentation** — dated inventories, control evidence, denial-of-access records — because political-violence and expropriation claims mature on evidence assembled early, and a durable loss matures into a claim. *Four:* a **preserved-title stance** — an occupier's paperwork is not legal transfer; title and arbitration claims survive contested ground and are long-horizon

instruments to hold, not near-term values to book; and no exit mechanics that could read as legitimizing an occupier's administration. *Five:* a **survival-map negotiation posture** — preserve licenses, relationships, and re-entry options with the surviving state, and treat occupied-territory assets as the legal channel's problem, not the going-concern channel's.

Package C — Pre-emptive restructuring. Commit capital now to the third shape: fire-sale the east, redomicile the network. Cost: irreversible moves on day 4, ahead of any confirmed geography, on one input.

The CFO's objection had real force: *"Russian troops are standing on Ukrainian soil this morning. 'They will hold some of it' isn't a forecast — it's a description of the map with extra steps."* And the honest half-concession Voss wrote into the minutes: the *loss* limb might indeed be momentum. What the visible front did not supply — what neither camp in the room held — was the conjunction: loss at scale, *and* a surviving state, *and* durability. The 72-hour school denied survival; the restoration school denied durability. The committee also logged what the statement could **not** buy: physical relocation. Movable property near a front that was shifting daily could not be extracted on the strength of a day-4 statement naming no location; the convertible value was balance-sheet, legal, and posture — all of it running on committee time, of which quarterly cycles gave them plenty, if the shape held.

EXHIBITS

Exhibit 1 — The two-camp consensus, week one. The "72-hour" collapse school versus eventual-restoration; calibrated partial-durable-loss neither camp's base case in the documented Western frame of the invasion's first week.

Exhibit 2 — The sealed input. Public analysis of February 28, 2022 (announced at x.com/vijayjyotish/status/1498384003543420928; re-posted in the open as a screenshot March 22, 2022): the three-limbed statement quoted verbatim, self-banded "appears very likely" (0.90 on the record's published ladder). Advisory documentation and grading at jyotishintelligence.com/intel/advisory/IA-RU-044 — including the record's claimed-versus-controlled discipline and its refusal to grade the untestable "without war" branch. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — Ostrava's exposure (composite teaching device). Two eastern plants; central distribution network; receivables and inventory; 300 employees (the gravest exposure, and the first workstream regardless of any map); PV/expropriation cover turning on characterization: temporary interruption vs. confiscation vs. durable deprivation.

Exhibit 4 — What subsequently happened (public record, for the teaching note). September 30, 2022 — 214 days after the seal: Russia formally annexes Donetsk, Luhansk, Kherson, and Zaporizhzhia oblasts — ~15% of Ukraine's recognized territory on paper, ~20% with Crimea — while Ukraine survives as a state. Claimed ≠ controlled throughout: Zaporizhzhia's capital never taken; Kherson city retaken November 2022; no broad international recognition. The de facto loss held and deepened — 2025 closed as Russia's largest territorial-gain year since the war's first. Every negotiation and reconstruction framework since has priced survival-with-durable-partial-loss.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Pick a binary	B · Segment on the third shape	C · Pre-emptive restructuring
What it adds	a clean model	geography as the unit; claims maturing on early evidence	finality
What it forecloses	the shape the war took	nothing	reversibility, on day 4
Cost if a binary proves right	0	a memo and some spare rigor	fire-sale discounts
Cost / value if the third shape holds	oscillation; mispriced everything	each tier priced to its actual path	right — by luck, at day-4 prices
Speed	—	one committee sitting	months, irreversibly
Authority required	committee's own	committee's own	board
Reversibility	revisable, expensively	full	poor
Recommendation & rationale		*(yours)*	
In hindsight		*(filled at grading)*	

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/third-shape
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Third Shape

Rails: the human gravity of the setting is acknowledged and never instrumentalized — staff safety is the first workstream under every package and no recovery arithmetic outranks it; the map-momentum objection is conceded in part (the *loss* limb) and the note must locate the divergence precisely in the conjunction (survival + durability), not in the loss; "significant" carries no numeric threshold and a stricter reader may discount it; the "without war" branch never executed and is graded by no one; physical relocation was **not** convertible on this seal; the grade is self-assigned under a published rubric — disclosed, not equivalent to external adjudication.

LEARNING OBJECTIVES

(1) The third-shape instinct: when a room holds exactly two clean maps, the highest-value analytic act is often to ask what shape *neither* camp is pricing — because binary frames make geography irrelevant and the excluded middle is where differentiated decisions live; (2) shape-first, geography-later: how to act on a structural claim that names no location — standing up the *register* now and populating it with judgment as facts arrive; (3) claims mature on early evidence: insurance characterization (interruption / confiscation / durable deprivation) as a decision made months before any adjuster asks, by what the desk documents now; (4) claimed-versus-controlled as a permanent discipline: paper transfers are not facts, de facto conditions are, and both the record and a recovery model must run on the latter; (5) conjunction analysis: decomposing a statement into limbs and asking which limb the visible world already supplied — the momentum objection's honest answer.

ASSIGNMENT QUESTIONS

1. Write the committee memo instantiating Package B's five items for Ostrava's actual exposure (Exhibit 3), with the register's two tiers defined operationally — who judges "elevated risk of durable loss," on what evidence, revised when?
2. The CFO: "It's a description of the map with extra steps." Decompose the sealed statement into its three limbs and assign each to *visible-front momentum* or *genuine divergence*, citing Exhibit 1. Which limb did the committee's decisions actually turn on?
3. Characterization drives coverage: for each of Ostrava's exposures, state which insurance characterization each of the three maps implies, and what Package B's documentation program must capture *now* for the durable-deprivation theory to be provable in 2024.

4. "Preserved title is a long-horizon instrument." Price it: what is a contested claim against a frozen line worth under each map, and what exit mechanics would destroy it?
5. Complete Exhibit 5 and sign it; fill the hindsight row as of October 2022 (annexation on paper) and again as of 2026 (the line frozen, the loss deepened). Does your recommendation survive both readings?
6. **The machine staff-officer.** An AI country-risk system ingests front-line mapping, consensus analyst notes, and this sealed statement. Trained on the room's two camps, it assigns the third shape a low prior — exactly the failure this case teaches. Design the prompt or process that forces a model to *enumerate excluded shapes* rather than interpolate between loud ones, and state which committee judgments (Exhibit 5 rows) remain human even when it works.

THE BOARD PLAN (90 MIN)

Open on Exhibit 3 and the two maps (15): run the committee's week under collapse and under restoration — students must feel how *clean* both binaries are, and that cleanliness is the trap. Introduce Exhibit 2 (15): Q2's limb decomposition on the board; give the CFO's objection full force and let the conjunction answer it — or fail to. Build Package B (20): the five items, each priced; insist on the register's operational definition (Q1's hard part). The insurance deep-dive (15): Q3 — this is the case's most transferable professional content. Commit (15): memo + Exhibit 5. Reveal Exhibit 4 (10): the 214-day arc, claimed-versus-controlled throughout; close on Q5's two-date hindsight — a recommendation that survives both readings was priced on the shape, not the news.

WHAT HAPPENED, AND THE GRADE

On September 30, 2022 — 214 days after the seal, two to three full committee cycles — Russia annexed four oblasts on paper (~15% of Ukraine's recognized territory; ~20% with Crimea) while Ukraine survived as a state; Kherson city was retaken that November, Zaporizhzhia's capital was never taken, and the de facto loss nonetheless held and deepened through 2025. The war settled into precisely the shape neither week-one camp had priced, and every negotiation framework since has run on it. The record grades the call a **HIT** at its self-banded 0.90 — and carries the concessions this note enforces: the loss limb was arguably momentum, "significant" is unquantified, the "without war" branch is graded by no one, and the adjudication is self-performed under a published, recomputable rubric. For the composite committee, the postures priced as the shape predicted: the binary desks oscillated — collapse's write-offs looked foolish by April, restoration's carrying values looked foolish by October — while the segmented register absorbed each month's map into tiers it had built in February, and the deprivation file, started early, matured into exactly the evidence its claims required.

EPILOGUE

The third-shape instinct is this case's export far beyond wartime: markets, elections, and technology transitions all generate two-camp rooms, and the excluded middle — *it survives, diminished, durably* — is history's most common outcome and analysis's least defended. The committee discipline is one question, asked every time the room converges on a clean binary: *what is the cheapest register we could stand up today for the shape nobody here is pricing?*

CASE TWENTY-TWO

The Festival Window: A Flag the Calendar Already Knew

A teaching case in public-safety operations — what a low-surprise, high-actionability warning is worth to a desk that already plans well

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the record's own calibration note concedes this call's central weakness up front: **"festival-window frictions are not rare,"** and no improbability figure was computed for it. The documented incidents are real and cited; the case treats communal violence with the gravity it demands, takes no side in any communal question, and **never decodes** the sealed claim's threat-class wording — the record carries it verbatim and uninterpreted, and so does this case. The **protagonist and his district are composite characters** — "Superintendent Arjun Nair" and "Chandrapura district" are not real; no actual police department's deliberations are depicted, and no outside-world fact discussed is invented. **No claim is made that any posture would have prevented any incident.** The case takes **no position** on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

SUNDAY, AUGUST 18, 2024 — 10:00, A DISTRICT HEADQUARTERS IN SOUTHERN INDIA

Superintendent Arjun Nair ran police operations for Chandrapura district, and his desk calendar already knew what September held: Ganesh Chaturthi, main day the 7th, followed by the visarjan season — immersion processions winding through mixed neighborhoods, the annual test of every bandobast plan in the state. His department had run that plan for decades. Festival windows were fixed high-alert dates; that was not intelligence, it was the calendar.

The item his special-branch inspector flagged that morning was, therefore, almost redundant — which was exactly what made it worth a teaching file. A sealed public forecast of **August 16** — part of an eleven-call national-security sitting from a source whose record was public, graded, and method-disclosed

as Vedic astrology — named the Ganesh Chaturthi window as requiring *"comprehensive security measures"* and *"heightened caution,"* flagged large gatherings as attractive to threats, noted that Mumbai had seen no major incident since the 2011 serial blasts and that it would be *"prudent not to underestimate the risk this year,"* and carried a threat-class phrase — "sleeper cells" — that the record itself refused to interpret, then or later.[1] No route, no district, no actor decoded. Twenty-six days before the peak of the season.

Nair's inspector put the objection on the table before Nair could: *"Sir, this tells us the sky is blue. Every SP in the country treats this window as high-alert. What is a flag worth when the calendar already carries it?"* And that, Nair decided, was the actually interesting question — not about astrology, but about **an entire class of input** his desk received constantly: advisories from higher formations, seasonal circulars, external tips — most of them naming risks his standing plan already covered. Some of those inputs deserved nothing. Some deserved to change the *quality* of execution of a plan that existed anyway. Which was which, and by what test?

THE THREE POSTURES

Package A — The calendar is enough. File the flag; run the standing bandobast plan exactly as every year. Cost: zero. Defensible fully — *if* the standing plan was actually being executed to its own standard rather than to its habit.

Package B — Use the flag as an execution audit. Four actions, all inside his own authority, all cheap, all things the standing plan nominally contained and a dated external flag gave him the occasion to actually verify, three weeks out instead of three days. *One:* route and force posture — walk the roster now; confirm patrol density, cordon placement, and camera coverage on the procession corridors as a planned reassignment rather than a scramble. *Two:* **prohibitory-order readiness** — pre-draft the Section 144 order, review the legal basis, pre-clear the sign-off chain, so the time from "clash begins" to "order in force" is minutes, not hours. *Three:* community liaison — hold the pre-festival meetings with community and religious leaders along every route *his own judgment* flagged as sensitive (the seal named none; the judgment was his department's to make, as it always was). *Four:* task the district intelligence unit to monitor procession-route chatter and tension indicators through the window — a tasking instruction to analysts already on payroll.

Package C — Surge. Treat the flag as specific warning: request additional companies from the range, elevate public messaging, visible saturation. Cost: real — force drawn from neighboring commitments, a public signal that could itself raise the temperature, and all of it resting on a single low-improbability call inside an eleven-call sitting.

What Nair wrote in the margin of the file was the test the inspector's question deserved: *the flag adds no information about whether the window is risky — the calendar owns that. What it might add is a dated occasion to convert a standing plan from paper to posture. If my Section 144 draft is current, my liaison meetings scheduled, and my rosters walked — the flag is worth nothing, and I should be proud of that. If any of those is stale, the flag just found it, and the finding cost me a morning.* The bandobast conference was Thursday.

EXHIBITS

Exhibit 1 — The standing doctrine. Ganesh Chaturthi as a fixed high-alert date in every observing state, backed by decades of institutional experience with procession-route friction; bandobast planning cycles typically compressing in the final days before the festival.

Exhibit 2 — The sealed input. Advisory IA-IN24-NS03, sealed August 16, 2024, within an eleven-call national-security sitting: the festival window named, the caution quoted verbatim in the case text, the threat-class phrase carried **undecoded** — by the record and by this case. The record's own honesty: **no improbability figure computed** ("festival-window frictions are not rare"); value claimed on actionability, not surprise; the sitting's full tally on the public record — **8 HIT, 2 NEAR, 1 MISS**, the miss retained at full weight. Documentation at jyotishintelligence.com/intel/advisory/IA-IN24-NS03. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The execution audit (composite teaching device). Standing-plan items nominally in place: route rosters (last walked: unknown); Section 144 draft (last reviewed: previous year); liaison calendar (unscheduled); intelligence tasking (routine, not window-specific). The gap between a plan that exists and a plan that is current.

Exhibit 4 — What the window held (public record, for the teaching note). The 2024 window produced a documented multi-state spike in procession-linked communal and security incidents — friction across more than a dozen states. Representative sourced flashpoint: Nagamangala, Mandya district, Karnataka,

September 11 — a clash as an immersion procession passed near a mosque; arson and stone-pelting; ~20 shops and 8 vehicles burned; 52 arrests; Section 144 imposed *after* the incident; a police inspector suspended over the handling. No fatality in the cited record. The record grades the call on the incident-spike class materializing — **not** on the undecoded threat clause, and not on Mumbai, which stayed quiet while the flashpoint landed hundreds of miles south.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Calendar is enough	B · Execution audit	C · Surge
What it adds	nothing — by design	a dated occasion to verify the standing plan	visible mass
What it forecloses	discovering staleness in time	nothing	neighbors' coverage; a calm signal
Cost if the window is quiet	0	a morning of verification	real force misallocated
Cost / value if friction lands	plan executes at its habit	plan executes at its standard	possibly helpful, possibly heat
Speed	—	days, within standing authority	range-level request cycle
Authority required	none	SP's own	range/state
Reversibility	n/a	full	politically sticky
Recommendation & rationale		*(yours)*	
In hindsight		*(filled at grading)*	

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/festival-window
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Festival Window

Rails: communal violence is treated with gravity and without side-taking; the threat-class phrase stays **undecoded** — any student attempt to interpret it is out of bounds, exactly as the record rules for itself; **no prevention claim** — nothing asserts that any posture anywhere would have changed Nagamangala, and the suspended inspector is a documented fact, not a lesson in blame; the standing-doctrine objection is conceded at full strength: for a department that already plans well, this flag's marginal value may be **close to zero**, and the case is built so that answer can win.

LEARNING OBJECTIVES

(1) The two value axes, separated: improbability (what case 18 scored) versus actionability — this flag is the mirror image of the Mamdani call, near-zero surprise but same-day actionable, and the SITA-style lens (60% of weight on impact and actionability) exists precisely because desks consume warnings with muscles, not with scoreboards; (2) the execution-audit use of redundant warnings: a flag that adds no information can still add an *occasion* — the cheapest possible trigger for converting a standing plan from paper to current; (3) tempo as the real gift: 26 days versus the compressed final-week cycle, priced in the difference between a pre-cleared prohibitory order and one drafted during a clash; (4) intake honesty at the bottom of the value range: how a record that computes no improbability number for its own call — and says why — models the discipline of not manufacturing surprise where none exists.

ASSIGNMENT QUESTIONS

1. Write Nair's Thursday bandobast-conference note adopting Package B — framed so that it neither over-credits the flag nor wastes the occasion, and survives a subordinate asking "since when do we audit rosters on an astrologer's say-so?"
2. The inspector's question, answered rigorously: define the test that sorts redundant-but-useful inputs from pure noise for a desk whose calendar already carries the risk. (Expected shape: the input's value = the staleness it can reveal × the cost of the occasion.)
3. The seal named Mumbai; the sourced flashpoint landed in Karnataka. What does this teach about consuming *place* words in warnings versus *class* words — and how should the record's own grading choice (incident-spike class, not geography) inform Nair's reading?

4. Compare this case to The City Hall Window (case 18): both inputs carry ~zero surprise. Why does one deserve a morning of the desk's time and the other deserves only a process fix? Derive the general rule.
5. Complete Exhibit 5 and sign it; fill the hindsight row as of September 12, honestly — remembering that Chandrapura is not Nagamangala, and a quiet district proves nothing in either direction.
6. **The machine staff-officer.** A state-level AI system now drafts bandobast execution audits automatically before every calendar-fixed festival window — rosters, order drafts, liaison schedules — regardless of any external flag. Does that make inputs like Exhibit 2 worthless? What did the flag actually supply (an occasion) that the machine now supplies on schedule — and what new staleness does the *machine's* checklist accumulate that only an outside prompt might catch? Write the standing order for auditing the auditor.

THE BOARD PLAN (75 MIN)

Open on Exhibit 1 and the inspector's objection (15): let "the calendar already knows" win the opening argument completely — the case only works if the room genuinely believes the flag is redundant. Introduce Exhibit 3 (15): the staleness inventory; watch the room rediscover that *knowing* a window is risky and being *current* for it are different states. The two-axes lens (15): Q4's comparison with case 18 on the board — surprise versus actionability as orthogonal axes, with both cases plotted. Commit (15): the conference note + Exhibit 5. Reveal Exhibit 4 (10): the multi-state spike, Nagamangala, the post-incident Section 144, the suspension — read plainly, no triumphalism; the record graded the class, not the geography, and Mumbai stayed quiet. Close (5) on Q2's test, stated as one line.

WHAT HAPPENED, AND THE GRADE

The window produced a documented multi-state spike in procession-linked incidents; the representative flashpoint at Nagamangala on September 11 — twenty-six days after the seal — brought arson, fifty-two arrests, a prohibitory order imposed *after* the clash, and an inspector's suspension. The record grades the call a **HIT on the incident-spike class only**, refuses the undecoded threat clause any credit, concedes the geography miss-match (Mumbai named, Karnataka burned), computes no improbability number because the base rate is high — and keeps the same sitting's outright MISS on the board beside its eight hits. The teaching synthesis is the volume's most modest and most practical: this flag never deserved belief, escalation, or a surge. It deserved a morning — the execution audit that a well-run desk would defend on its own arithmetic — and the honest verdict that for a department already current, the correct price was zero. Knowing which of those two departments yours actually is, three weeks early, is what the morning buys.

EPILOGUE

Every profession has its festival windows — the quarter-close, the renewal season, the launch week — dates the calendar already marks and plans already nominally cover. The redundant warning is the most common input any desk receives, and this case's rule travels: *its value is never information; it is the staleness it can surface, priced against the cost of looking.* A desk that can extract that value from a source it finds absurd — and can also say "zero, we're current" and mean it — has an intake process, not a filing habit.

CASE TWENTY-THREE

The Do-Not-Spend Decision: Closing a Line and Holding It Closed

A teaching case in resource allocation, commitment devices, and why boring calls are the audit of interesting ones

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the record's own file rates this **the highest-prior, least-informative call of its series**, holds its contrarian score at 1-of-5, claims no premium, and concedes that any competent analyst would have written the identical memo for free from public polling. Every election figure and forecaster rating is real and cited. The **protagonist and his campaign are composite characters** — "Marcus Webb" and his budget committee are teaching devices; no real campaign's internal deliberations are depicted, and no outside-world fact discussed is invented. The case is strictly **nonpartisan**: the 2024 cycle is treated as an allocation environment, and nothing herein is political advice or comment.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

THURSDAY, AUGUST 29, 2024 — 11:00, A NATIONAL CAMPAIGN HEADQUARTERS

Marcus Webb ran resource allocation for a national presidential campaign, sixty-eight days from the vote, and his job description fit in one sentence: *every dollar spent where the race isn't is a dollar unavailable where it is.* The hard half of the job was not finding money for the contested seven. It was refusing money to everywhere else — because a national campaign's default posture toward every state is "monitor," and monitoring is cheap to authorize, expensive to defend *not* doing, and never, ever, voluntarily concluded. Tracking polls refresh. Field-office scouting trips get penciled. Ad-buy options are "held in reserve." A state line, once open, stays open by institutional gravity.

Minnesota was Webb's cleanest test case. Ten electoral votes; Democratic every cycle since 1976; the sitting governor now on the opposing ticket as of twenty-three days ago; every named public forecaster — Silver's model, Cook Political Report ("Solid Democratic"), the dominant in-state KSTP/SurveyUSA poll — carrying it safe for the opposing side with no tipping-point scenario touching it. His own pollsters said the same. And yet the Minnesota line was still open: a tracking refresh scheduled for mid-September, a "reserve" ad-buy option, a scouting memo in someone's drafts — because closing a line requires a decision, and monitoring requires only inertia.

The input that put the decision on Thursday's budget-committee agenda was, informationally, nothing. A public post of that morning, 8:21 AM ET — ninth in an eleven-day, state-by-state sealed run from a source whose record was public, graded, and method-disclosed as Vedic astrology — called Minnesota outright for the opposing ticket and added, as operational advice: *"it would be prudent… to refrain from allocating resources to this state, as the likelihood of altering its outcome is minimal."*[1] Webb's deputy summarized it fairly: *"An astrologer agrees with Cook, Silver, our pollsters, and forty-eight years of history. So?"*

The *so*, Webb thought, was not Minnesota. It was October. He had run cycles before, and he knew what the closing weeks did to closed lines: someone senior floats the safe state as a "stretch target," an expansion narrative wants feeding, and the line he zeroed in August gets re-litigated in the exact weeks when a redirected dollar was worth most and analysis was worth least. What the strange post supplied — what all four of his converging sources actually supplied, if he bundled them — was the material for a **commitment device**: a written, dated, falsifiable closure memo, signed now, designed to be *harder to overturn in October than it was to write in August*.

THE THREE POSTURES

Package A — Monitor. Leave the line open at its trickle: the tracking refresh, the reserve option. Cost: small in dollars, large in principle — the discovery pipeline keeps consuming attention, and the line remains available for October's stretch-target fight, unarmored.

Package B — Zero-spend closure. A same-week budget-committee decision, inside his existing authority: zero the Minnesota discovery budget — no tracking refresh, no scouting, no reserve option — with the freed dollars and staff-weeks pre-assigned to the contested seven; and the closure memo

written as a commitment device: the four convergent sources cited, the date fixed, and a standing clause that *any argument to reopen must defeat this memo's August logic, not merely restate October's mood.*

Package C — Confirm and close. Buy one more confirmatory in-state poll, run the full sign-off chain, close with ceremony. Cost: real money and weeks spent confirming a line no named forecaster had ever contested — the monitoring reflex wearing a decisive costume.

The deputy's real objection came at the memo draft: *"Why does the astrologer appear in it at all? Cook and Silver are the citation; the other thing is a liability if it leaks."* Webb's answer was the case's quiet center: the post earned its line in the memo not as evidence about Minnesota but as **convergence across independently produced reads** — his pollsters, the public models, and a source whose method shared nothing with either had all closed the same state; and convergence across *differently generated* sources, he argued, is precisely the signal that a line is safe to close early rather than merely probably safe. Whether that argument survived the leak test was the last open question before the committee sat.

EXHIBITS

Exhibit 1 — The board, August 29, 2024. Minnesota: Democratic since 1976; home-state governor on the opposing ticket since August 6; Silver final-model safe, Cook "Solid Democratic" throughout, KSTP/SurveyUSA never showing a lead for Webb's side. Campaign default: every state line open at "monitor."

Exhibit 2 — The sealed input. Public post of August 29, 2024, 8:21 ET (x.com/vijayjyotishusa/status/1829177597055644083): the state called outright; the refrain-from-allocating advice quoted verbatim. The record's own pricing at jyotishintelligence.com/intel/advisory/IA-US24-010: floor 1-in-4 (~2 bits — near the bottom of its corpus), contrarian score 1-of-5, **no premium claimed**; the same nine-state run carries companion calls graded **MISS** (Georgia, Michigan, Nevada) on the same public ledger, retained at full weight. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The gravity of "monitor" (composite teaching device). Open line contents: scheduled tracking refresh (mid-Sept); reserve ad-buy option; scouting memo in drafts. Known October dynamics: stretch-target narratives, expansion pressure, senior-principal enthusiasm — historically strongest in the exact weeks redirected dollars are worth most.

Exhibit 4 — What subsequently happened (public record, for the teaching note). Minnesota certified for Harris by 4.2 points (50.97–46.75, ~185,000 votes), all ten electors — tighter than 2020's 7-point margin, but never in reach. The public record shows the Republican campaign left the state without a statewide ad buy or ground program (one rally, in July, pre-seal); closing-weeks campaign chatter floated Minnesota as a "stretch target" — the reopen pressure the sealed copy had pre-named as miscalibrated. The contested seven got the infrastructure.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · Zero-spend closure	C · Confirm and close
What it adds	optionality (mostly illusory)	freed dollars + an armored NO	ceremony
What it forecloses	nothing, ever — that's the problem	the October re-litigation	weeks + a poll's cost
Cost if the state tightens for real	trickle continues	the memo's falsifier triggers reopening honestly	same as B, later
Cost / value in the base case	attention drained all fall	marginal dollars at the margin that matters	B minus the wasted confirmation
Speed	—	same week	2–3 weeks
Authority required	none	the director's own	polling budget sign-off
Reversibility	n/a	by defeating the memo — deliberately hard	same
Recommendation & rationale		*(yours)*	
In hindsight		*(filled at grading)*	

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/do-not-spend
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Do-Not-Spend Decision

Rails: strictly nonpartisan — an allocation environment, not a political argument; the honest-weakness concession leads every discussion: this was the series' highest-prior call, worth ~2 bits, and **any competent analyst could write the identical memo for free** — the case's subject is the discipline around such calls, never their difficulty; no claim that the sealed post influenced any real campaign's budget (no internal minutes are public); the GA/MI/NV companion **misses** stay visible — they are the reason the case exists.

LEARNING OBJECTIVES

(1) Closing versus monitoring: why "keep watching" is never neutral — the open line's true cost is attention, optionality theater, and its availability to October's worst arguments; (2) commitment devices: designing a decision to be harder to reverse casually than it was to make carefully — the dated memo with a falsifier clause as the allocation-desk version of the documented hold; (3) convergence across independently generated sources as the close-early signal — and its leak-test limits; (4) **the ledger-completeness principle**, this case's reason for being: a record that ships its boring, over-determined calls on the same clock, hash, and grading rule as its contested ones — and keeps its misses beside them — is the only kind whose contested calls a director can rationally weight; Minnesota is the *audit*, not the discovery.

ASSIGNMENT QUESTIONS

1. Write Webb's closure memo: the convergent citations, the falsifier clause (what evidence *would* legitimately reopen the line?), and the October-proofing language. Decide — and defend — whether the sealed post appears in it.
2. The deputy's leak objection versus Webb's convergence argument: adjudicate. Is a method-independent source worth extra convergence weight, zero, or negative (contamination risk)? State your rule for unconventional sources in *internal* citations.
3. Compute the case's own arithmetic: if monitoring Minnesota costs X in polls and staff-weeks and the contested-seven marginal dollar returns Y, what X/Y makes Package A defensible? Now add October: price the *expected* cost of an unarmored line meeting a stretch-target narrative.
4. The ledger-completeness principle, applied: you are a diligence analyst evaluating this record's *contested* 2024 calls. Walk exactly how the Minnesota call and the GA/MI/NV misses — together — change what the hits are worth. What would their

absence have told you?

5. Complete Exhibit 5 and sign it; fill the hindsight row as of certification, remembering the margin tightened by nearly three points and the electors never moved. Does tightening-without-flipping vindicate or indict the closure?
6. **The machine staff-officer.** An AI allocation model rates all fifty states continuously and would have closed Minnesota in July. It would also — optimizing expected electoral votes — quietly reopen it if a single outlier poll landed in October. What does the commitment-device logic teach about *when a desk should want its optimizer to be sticky*, and which Exhibit 5 rows encode that stickiness? Write the override protocol.

THE BOARD PLAN (75 MIN)

Open on the gravity of "monitor" (15): Exhibit 3; have the class list what an open line consumes — the insight that monitoring is a *decision with costs* must be theirs before any forecast appears. The input lands (10): let the deputy's "So?" win; the call is worth ~2 bits and the room should say so. The October turn (15): introduce the stretch-target dynamics and re-ask what an August memo is *for* — objective 2 emerges here. The convergence fight (15): Q2 in pairs, leak test included. Commit (10): memo + Exhibit 5. The ledger-completeness close (10): Q4 walked carefully — this is the case's deepest content and the volume's own epistemics made explicit: the boring call and the kept misses are why anything else on the record is legible.

WHAT HAPPENED, AND THE GRADE

Minnesota certified for Harris by 4.2 points; all ten electors held; the tightening from 2020 never put the state in reach. The public record shows the Republican side ran no statewide buy or ground program there — and that closing-weeks chatter floated the state as a stretch target anyway, the exact reopen pressure a commitment memo exists to meet. The record grades the call a **HIT priced at almost nothing**: floor 1-in-4, ~2 bits, contrarian score 1-of-5, no premium claimed — shipped on the same clock and grading rule as the same run's Georgia, Michigan, and Nevada calls, which resolved **MISS** and remain on the ledger at full weight. That pairing is the grade that matters. A record that only surfaced its contested wins would be an advertisement; one that seals its over-determined calls and keeps its losses is an instrument a diligence desk can actually calibrate against — which is what a resource-allocation director is, every day, for every source in the building.

EPILOGUE

The volume's readers will notice this case rhymes with The City Hall Window: both inputs are worth ~2 bits, and both cases locate the value somewhere other than the call. The muni desk learned to score sources on surprise; the allocation desk learns the complement — that *unsurprising calls, honestly sealed, are the audit trail that makes*

surprising ones bankable. Together they are this book's answer to the question every skeptic opens with: a record is not its highlights. It is its clock, its denominators, and its refusal to let the boring and the broken disappear.

CASE TWENTY-FOUR

The Long File: Designing for the World That Will Exist

A teaching case in policy architecture — the design premise as the cheapest decision in the largest instrument ever built

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the case carries the record's own concessions throughout: the headline "sanctions will fail" call may simply be **the base-rate position** (sanctions scholarship has long found maximalist objectives rarely achieved); the sealed language is directional and unquantified; and the call remains **OPEN through 2035** — its grade is a tracking judgment, not a closed verdict. Every economic figure, institutional forecast, and public statement is real and cited. The **protagonist and her unit are composite characters** — "Clara Lindqvist" and her drafting cell are teaching devices; no real ministry's deliberations are depicted, and no outside-world fact discussed is invented. The case takes **no position** on sanctions policy itself, and none on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

MONDAY, FEBRUARY 28, 2022 — 08:00, A COALITION CAPITAL

Clara Lindqvist drafted sanctions architecture — the working level where measures become instruments — and the week behind her had produced the fastest, largest coordinated economic-statecraft program ever assembled: SWIFT disconnections, a central-bank asset freeze, export controls, with much more coming. The week ahead would begin designing everything the first wave hadn't touched — above all, the energy-trade instrument that would eventually become the G7 oil price cap. And every design meeting she attended ran on the same premise, spoken and unspoken: **the Russian economy is going to collapse.** The IMF would soon project −8.5% for the year; the IIF said −15%; Goldman −10%; the French finance minister promised "the collapse of the Russian economy"; within a month the American president would say it was "on track to be cut in half."

The document her intake colleague flagged that morning said the premise was wrong. A public post of **February 28** — day four of the invasion, timestamped, from a source with a sealed record whose disclosed method was Vedic astrology — assessed the program in the words preserved at the source: *"The international sanctions appear unlikely to succeed in realizing their strategic objective of punitive consequences for Russia. It is most likely to be a paper-tiger than anything of real consequence."*[1] (A restatement that November would name the mechanism outright: trade rerouting to **India and China**, the treasury sustained by **oil, coal, and forest exports** — and an April 2023 addendum would seal a ruble arc of volatility followed by central-bank-restored stability.)

Lindqvist knew the honest first objection herself, because she had read the literature the collapse camp apparently hadn't: sanctions **rarely** achieve regime-breaking objectives — that was the field's own long-standing finding, and on the headline, the strange post might just be restating the base rate against a moment of political fever. But the drafting question in front of her was not "who wins the forecast." It was a **design premise**, and premises are the cheapest and most consequential objects in her trade: an instrument drafted on *collapse-and-marginal-substitution* prices its enforcement architecture, its third-country posture, and its designation resourcing one way; an instrument drafted on *full-scale substitution through non-coalition corridors from day one* prices all three the other way — and the premise could be changed at the speed of a modelling meeting, for free, this week, while every downstream decision was still open. The crude price cap's design window would run for months; whichever premise entered it first would shape everything the coalition built.

THE THREE POSTURES

Package A — Monitor. Keep the collapse premise; note the dissent. Cost: zero today. Exposure: if substitution ran at scale, every enforcement mechanism would be discovered *after* its corridor matured — retrofitted, package by package, against relationships already formed.

Package B — Substitution-first design. Change the premise at drafting: assume full-scale rerouting through the non-coalition corridors and treasury channels the input class named, and let the first draft — not the fourth amendment — price the designation strategy, the corridor monitoring, and

the intermediation posture. Cost: nothing. It is a modelling premise, inside existing authority, requiring no coalition agreement — the measures still get drafted; they simply get drafted for the world in which the target adapts.

Package C — Retrofit-anticipating overbuild. Hedge every conceivable evasion path in the first instrument. Cost: scarce drafting time spread across imagined corridors, the package delayed, the coalition strained — over-insurance against a single unconfirmed input.

Her section head's challenge was the room's fever, stated politely: *"Every serious institution has a number, and every number says collapse. You want the instrument designed around a Telegram post?"* Lindqvist's draft reply carried the case's whole argument: *No — around a premise. The post is one input; the sanctions literature is another; and both point the same way against five institutions who are, this week, marking to politics. If they're right, a substitution-first instrument wastes nothing — enforcement designed for adaptation still works on a collapsing economy. If we're right, a collapse-first instrument spends its first two years chasing corridors it assumed away. The premise is free. The retrofit is not.* The design meeting for the energy instrument was Thursday.

EXHIBITS

Exhibit 1 — The consensus, week one. IMF (April 2022 WEO): −8.5% Russian GDP for 2022; IIF −15%; Goldman Sachs −10%; Le Maire (Mar 1): "we will bring about the collapse of the Russian economy"; Biden (Mar 26): economy "on track to be cut in half." The sanctions-scholarship base rate, standing quietly on the other side: maximalist objectives are rarely achieved.

Exhibit 2 — The sealed input. Public post of February 28, 2022 (Telegram, announced on X; advisory IA-RU-001 at jyotishintelligence.com/intel/advisory/IA-RU-001): the paper-tiger assessment quoted verbatim; restated November 19, 2022 with the mechanism named (India/China rerouting; oil, coal, forest exports) — sixteen days before the crude cap took effect; ruble arc addendum April 30, 2023. Sealed 792 days before the materialisation milestone; **OPEN through 2035** — graded as a tracking judgment, not a closed call. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The design windows (public record). Seal → crude price cap in force: 280 days (Dec 5, 2022); products cap: Feb 2023; the November restatement → crude cap: 16 days — the last moment the corridor assumption

could enter the original design rather than the retrofit. Authorities: Commission drafts/Council adopts; Price Cap Coalition sets design; OFAC/OFSI designate in days-to-weeks under existing statute.

Exhibit 4 — What subsequently happened (public record, for the teaching note). Russian GDP 2022: −2.1% (a seventh of the worst institutional forecast); 2023: +3.6%; 2024: +4.1% — faster than every G7 economy (IMF, after five consecutive upward revisions). India–Russia trade: record $66B (FY2023–24), Russia at 40%+ of Indian crude imports; China–Russia trade: $240B+ (2023). Oil revenues 2023: $103B; budget surplus restored 2024; real wages +8.7%. Ruble: mid-2023 volatility, emergency 350bp hike (Aug 15, 2023), stabilization behind capital controls. GAO: sanctions "may have had a smaller effect than ... previously estimated"; Fed Chair Powell, candid on a recorded prank call: Russia has "been able to avoid the harshest application... they're doing everything they can to sell their oil to the Chinese and the Indians, and they're buying it." By mid-2024, "sanctions have failed" framing ran in Bloomberg, the FT, and the Economist — roughly two years behind the seal. The enforcement programme retrofitted across 2023–2025: shadow-fleet designations, intermediary-refiner pressure, guidance waves — each built after its corridor matured.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Monitor (collapse premise)	B · Substitution-first design	C · Overbuild
What it adds	consensus comfort	corridors contested while thin	imagined completeness
What it forecloses	concurrent design	nothing — collapse-world still served	drafting tempo; coalition patience
Cost if collapse comes	0	~0 (adaptation-ready design still works)	delay, for nothing
Cost / value if substitution runs	two-year retrofit, corridors hardened	enforcement born where the trade goes	partially right, exhausted
Speed	—	one modelling meeting	months
Authority required	none	drafting cell's own	coalition-wide
Reversibility	expensive (the retrofit is the reversal)	full	poor
Recommendation & rationale		*(yours)*	

In hindsight		*(filled at grading — and this file stays OPEN)*	

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/long-file
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Long File

Rails: the base-rate concession leads — a hostile economist reads the headline call as the literature's own position, not clairvoyance, and the case's design argument must therefore run on the **mechanism-level content** (named corridors, named treasury channels), which no named institution had on the record; the sealed language is directional and unquantified; the call is **OPEN through 2035** and its grade is a tracking judgment — the class should feel what it means to teach from an unclosed file; substitution had willing sovereign counterparties, and **no instrument design forces India or China to decline discounted crude** — the counterfactual's ceiling is enforcement cost and tempo, never the macro outcome; the case takes no position on sanctions policy.

LEARNING OBJECTIVES

(1) Premises as the cheapest decisions: in any large instrument — legislation, architecture, procurement — the modelling premise is set early, for free, and silently prices everything downstream; retrofit is the expensive spelling of "wrong premise"; (2) separating headline from mechanism: the failure call was arguably base-rate, the *named corridors* were not — and only the mechanism content was designable-against (the volume's decomposition discipline, at policy scale); (3) fever versus literature: recognizing when institutional consensus is marking to politics against its own field's base rates — and what a dated outside input contributes then (cover, a falsifiable artifact, an occasion — not information the literature lacked); (4) open files: how to weight, cite, and teach a call that has not finished resolving — the tracking judgment as a permanent state of most real policy evaluation.

ASSIGNMENT QUESTIONS

1. Write Lindqvist's Thursday design memo adopting Package B — the premise stated, the base-rate literature cited first, the sealed input cited honestly or omitted deliberately (defend the choice), and the falsifier named: what evidence by mid-2023 would send the premise back?
2. Decompose the input: which of its limbs (failure headline / corridor names / treasury channels / ruble arc) was base-rate, which was mechanism, which was designable-against? Map each limb to a concrete drafting difference in the energy instrument.
3. "An adaptation-ready design still works on a collapsing economy." Test that asymmetry claim rigorously — construct the scenario where Package B *does* cost something real, and price it against the documented retrofit.

4. The two-year lag between the seal and the institutional reversal is Exhibit 4's quiet spine. Using the tempo entries in Exhibit 3, identify the three moments where the corridor assumption could have entered the design, and what was lost at each pass.
5. Complete Exhibit 5 and sign it — then write the hindsight row *as a tracking judgment*, dated today, knowing the file stays open to 2035. What would move your entry, in either direction?
6. **The machine staff-officer.** A coalition AI now war-games every draft instrument against simulated adaptation — reroutes, shadow fleets, parallel payments — before adoption, making "substitution-first" the automatic premise. What does the two-year human lag of 2022–24 teach about *why* the collapse premise survived so long (incentives, not information) — and which part of that failure mode does the machine inherit rather than fix? Write the paragraph for the AI system's design charter that addresses it.

THE BOARD PLAN (90 MIN)

Open on the instrument (15): what a sanctions package *is* — measures, premises, enforcement architecture — and where premises live; students must see that the premise is a decision before they see anyone contest it. The consensus fever (10): Exhibit 1, with the literature's base rate read against the ministers' quotes — the room should notice the institutions disagreeing with their own field. The input and the decomposition (20): Q2 on the board, limb by limb; the mechanism content must emerge as the only designable part. The asymmetry fight (15): Q3 — Package B's "free" claim deserves adversarial pressure. Commit (15): memo + Exhibit 5, with the falsifier clause mandatory. Reveal Exhibit 4 (10): the inversion, the concessions, the Powell call; run Q4's three-moments timeline. Close (5) on the open file: some verdicts don't arrive — desks decide anyway.

WHAT HAPPENED, AND THE GRADE — SO FAR

The consensus inverted. Russia contracted 2.1% where institutions had said 8.5–15; grew faster than every G7 economy in 2023 and 2024; ran record trade through precisely the named corridors; restored its budget surplus behind the named treasury channels; and rode the sealed ruble arc through an emergency hike to stability. The coalition's own institutions conceded in their own words — the GAO in a report, the Fed Chair on a recorded prank call naming the India/China reroute exactly. The enforcement programme was retrofitted across 2023–2025, corridor by matured corridor, at a pace the record's tempo entries measure precisely: the mechanism restatement preceded the crude cap by sixteen days, and the institutional reversal by two years. The record carries the call at 792 days to its materialisation milestone and holds the file **OPEN through 2035** — a tracking judgment, with the base-rate concession printed beside it — and this note grades the classroom's real subject the same way: not "was the astrologer right about sanctions," but whether a drafting cell, offered a free premise change on day one, with the field's own literature pointing the same direction, needed anything more than an occasion to take it.

EPILOGUE

This is the volume's longest file and its last, and it closes the arc the first case opened: a procurement office pricing a launch day, and a coalition pricing the largest instrument in economic history, face the same three-line discipline at opposite scales — *what does this input make cheap; what premise does my current plan silently carry; and what would it cost to design for the world that will exist instead of the one the room is announcing?* The record's answer, across twenty-four cases, has never been "believe the forecast." It has been: date the premise, price the alternatives, sign the package, and let reality referee — with the misses kept, the free wins refused, and the file open as long as the world is.

CASE TWENTY-FIVE

The Warning Delivered Twice: A Conversion-Cost Function

A teaching case in counterterrorism intake — what a broad, real warning costs to convert, and the tier at which breadth stops dividing value

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record. This case concerns, in its teaching note, a real atrocity in which more than 145 people were murdered; it is written with reverence for them, and **nothing in it claims that any triage decision, anywhere, would have prevented what happened** — the record itself states that counterfactual is unprovable and does not claim it. The **protagonist, her centre, and the warning text she receives are teaching devices**: "Nadia Osei" and the "Metropolitan Threat Integration Centre" are not real, and the warning in the case body is a declared structural transposition of the real sealed warning — identical in shape (a named city, a two-month window, a target class, a psychological objective; no venue, date, or perpetrator) — so that the class can price the decision without this case depicting any real agency's deliberations. Every fact in the teaching note's documented chain is real and cited. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

MONDAY, SEPTEMBER 11, 2023 — 07:30, A METROPOLITAN FUSION CENTRE

Nadia Osei had run the Metropolitan Threat Integration Centre for six years, and she could describe her hardest professional problem in one sentence: *most weeks bring a warning that is too broad to act on and too specific to file.* The tips line, the liaison cables, the walk-ins, the open-source flags — the stream never carried convenient warnings. It carried shapes like the one at the top of this morning's triage stack, submitted by her open-source cell from a publicly sealed, timestamped record whose disclosed method — she read the line twice — was Vedic astrology:

Terrorist attacks on public gatherings are possible in [the city] in March and April, in an attempt to sow panic among a public that has been insulated from hard realities.

Four elements, each load-bearing: a **city**. A **two-month window**, six months out. A **target class** — public gatherings. A **psychological objective** — panic as the point. And three absences, each equally load-bearing: no venue, no date, no perpetrator. Her deputy's triage annotation was the annotation every fusion centre writes on every warning of this shape: *"Unactionable as stated. What are we supposed to do — harden every concert hall in the city for sixty days?"*

Osei had learned to hear that sentence differently than her deputy meant it, because it contained an unexamined assumption about *price*. The deputy was pricing the warning at the only tier he could imagine acting at — venue-by-venue hardening — and at that tier he was right: the warning's honesty about what it didn't know (which venue, which day) divided its value across every candidate site, and the division drove the per-site value toward zero. But that arithmetic was a fact about the *tier*, not about the *warning*. The question she wrote in the margin was the one this case exists to teach: **which tier am I pricing this at?**

THE THREE TIERS

Tier 0 — Distribute. Route the item to liaison desks and partner agencies as an information product, no operational claim attached. Near-free; hours. Also, she noted, nearly meaningless on its own: distribution moves a warning around; it converts nothing. A warning can be distributed perfectly, repeatedly, to exactly the right offices — and still never become a single changed posture anywhere. Distribution *feels* like action precisely because it is the tier at which no one has to pay.

Tier 1 — The generic posture bulletin. One instruction, issued once, to a class of recipients, for a bounded period: a city-wide advisory to public-gathering venue operators — concert halls, stadiums, transit hubs — requesting elevated screening posture and a review of walk-in security through the named window. A request, not a mandate; staff-hours, not capital; reversible the day the window closes. And the property that made it the crux of her margin note: **its cost does not scale with the number of venues.** The same breadth that made the warning worthless at Tier 2 — *any* public gathering — is what a Tier-1 bulletin covers in a single act, for one fixed price.

Tier 2 — Venue-by-venue hardening. Named sites, armed screening, access-control changes, procurement, staffing. Cost scales with the candidate-venue count multiplied by the window's width. This was the tier her deputy's annotation had silently demanded the warning satisfy before crediting it — and the tier no sixty-day, no-venue warning can cheaply satisfy, which is why warnings of this shape are refused at the door of fusion centres everywhere: the chief is being asked, in effect, to fund a search of every candidate before acting on any.

The conversion-cost function, written for the training file she kept for her successors: *let W be the window's width and V the number of plausible venues in the named class. Tier 2 costs scale with V × W — breadth divides value. Tier 1 costs are invariant in V — breadth, at this tier, is coverage. Breadth divides a warning's value only at the tier the chief chooses to pay for.* Against Tier 1's modest price stood the real objections: her centre's alarm credibility with venue operators was a spendable currency, sixty days is a long posture, and the source was — she did not soften it in her own file — an astrologer, whose record her open-source cell described as sealed, graded, and miss-keeping, and whose method she had no framework to evaluate. The weekly triage board met Thursday. The item needed a tier.

EXHIBITS

Exhibit 1 — The triage economics (composite teaching device). Weekly intake: dozens of items; the broad-real shape ("city + window + class + intent, no venue/date/actor") recurring constantly from every source class — tips, liaison, open source. Standing failure mode: items priced only at Tier 2, found unconvertible, filed. Alarm credibility: finite; every bulletin spends some.

Exhibit 2 — The real sealed input (the shape's documented instance). Sealed September 4, 2023, on the public record (x.com/vijayjyotish/status/1698760248033476914): *"Terrorist attacks on public gatherings are possible in Moscow in March and April in an attempt to sow panic amongst the general public of Russia, who have been well-insulated from the harsh realities of this conflict."* City, two-month window, target class, psychological objective; no venue, date, or perpetrator. Advisory documentation, grading, and the full delivery chain at jyotishintelligence.com/intel/advisory/IA-RU-008 (Information Yield 8.71 bits at a published ~1-in-420; SITA 72 — specificity 86, impact 100, and **actionability 54, the record's own honest number for**

exactly the conversion problem this case prices). Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The consensus at seal time. September 2023's dominant counterterrorism lens inside the Russia–Ukraine theatre: the kinetic front. ISIS-K external-operations reach widely assessed as degraded. No elevated public threat baseline for Moscow gatherings. The warning bet against silence.

Exhibit 4 — The Decision Package, priced (recommendation left to the student).

	A · Tier 0: Distribute	B · Tier 1: Posture bulletin	C · Tier 2: Venue hardening
What it adds	the warning moves	one instruction covers every candidate venue at once	named-site protection
What it forecloses	nothing — and converts nothing	some alarm credibility, spent deliberately	budget, staffing, the next warning's hearing
Cost if the window passes empty	0	a bulletin recalled; credibility drawn down slightly	V × W — enormous, visible
Cost / value if the class is real	delivery without conversion	elevated posture at every gathering, for one fixed price	protection where you guessed; nothing where you didn't
Speed	hours	days	weeks–months
Authority required	none	the chief's own	city + venue owners + budget
Reversibility	n/a	full at window's end	poor
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/warning-delivered-twice
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Warning Delivered Twice

Rails — read aloud, and hold the room to them: more than 145 people were murdered at Crocus City Hall, and this note treats them with reverence — no triumphalism about the warning, no scoring language near the toll; **no prevention claim exists anywhere in this case** — the record itself states the counterfactual is unprovable and refuses it; the political-environment confound is the strongest alternative reading and gets full weight: the documented dismissal occurred inside a government with a standing incentive to characterize Western-sourced warnings as provocation, and a fusion centre without that incentive may not exhibit the failure; the case body's warning is a **declared transposition** — the real warning named Moscow, and no real agency's internal deliberations are depicted anywhere.

LEARNING OBJECTIVES

(1) The conversion-cost function: breadth divides a warning's value only at the tier the reader demands — Tier-2 costs scale with V × W while Tier-1 costs are breadth-invariant, so the honest question about a broad-real warning is never "is it specific enough?" but "which tier am I pricing it at?"; (2) distribution is not conversion: a warning can reach the right office twice and change nothing — intake processes must be audited on postures changed, not products routed; (3) the Two Warnings discipline: "is it real?" and "can I use it?" are separate questions with separate grades, and the record's own actionability score (54 — its lowest axis) models the honesty of keeping them apart; (4) alarm credibility as the Tier-1 budget constraint: the real price of the bulletin, and why it must be spent deliberately rather than hoarded to zero.

ASSIGNMENT QUESTIONS

1. Write Osei's Thursday tier decision and the bulletin's first paragraph — including how (or whether) the source is characterized to venue operators, and the recall condition when the window closes.
2. Formalize the conversion-cost function for your own city: estimate V for "public gatherings," price Tier 2 honestly, and find the crossover — how confident would a warning have to be before Tier 2 clears? What does that number teach about why Tier 2 is almost never the right demand?
3. The deputy's annotation — "harden every concert hall for sixty days?" — commits a specific analytical error. Name it precisely, and rewrite the annotation as it should read under the tier framework.

4. The teaching note's documented chain shows a warning delivered twice to a responsible office, with no documented conversion, and public dismissal three days before the event. Locate the failure *between* tiers — then apply the confound honestly: how much of this failure travels to a centre without the dismissal incentive, and how would you test whether yours has an equivalent incentive wearing different clothes?
5. Complete Exhibit 4 and sign it. Then write the hindsight row for both worlds — the empty window and the other one — and notice which entry is harder to write honestly.
6. **The machine staff-officer.** An AI triage system can price every incoming warning against the conversion-cost function automatically — V estimated from city data, tiers costed, bulletins drafted. It can also, trained on decades of filed warnings, learn that the base rate rewards filing. What does the documented chain teach about where warning systems actually fail (conversion, not detection) — and write the one metric a chief should hold an AI triage layer to, such that "distributed perfectly, converted never" scores as the failure it is.

THE BOARD PLAN (90 MIN — SCHEDULE THIS CASE WITH CARE; IT CARRIES REAL WEIGHT)

Open on Exhibit 1 and the deputy's annotation (15): let the room agree with him — the case works only if "unactionable" feels obviously true first. Build the three tiers and derive the conversion-cost function on the board (20): Q2's arithmetic, done live; the crossover number is the session's pivot. Q3 (10): rename the error. The commitment (15): Q1 and Exhibit 4, signed before any outcome is known. Then the teaching note's chain (20), read slowly and plainly — seal at D-200; hand-delivered in public to the office positioned to act at D-80; the U.S. Embassy's own alert at D-15; paired in real time at D-8; hand-delivered again at D-5; publicly dismissed as "blackmail" at D-3; the attack at D-0; more than 145 dead; ISIS-K's claim within hours; the after-action verdict of a serious triage failure. No editorializing — the ladder speaks. Close (10) with Q4's confound, given full honest force, and the question the class should carry out of the room: *what does my own intake convert — and what does it merely move?*

WHAT HAPPENED, AND THE GRADE

On March 22, 2024 — 200 days after the seal — four ISIS-K gunmen attacked the Crocus City Hall concert venue on Moscow's western edge, killing more than 145 people and injuring over 600: Russia's deadliest terrorist attack since Beslan. Every axis the sealed text had named resolved against it — the city, the window (inside March), the target class, and the intent, which the Russian government's own acknowledgment confirmed. The documented chain between seal and attack is the case's real subject: the warning accumulated over 650,000 impressions; it was posted, in public, to the account of the Deputy Chairman of Russia's Security Council on January 2 and again on March 17; the U.S. Embassy issued its own imminent-attack alert on March 7; and on March 19 the FSB Collegium was told Western warnings "resemble outright

blackmail." The record grades the call a **HIT** at 8.71 bits against a consensus that had ISIS-K's reach degraded and Moscow's threat baseline flat — and holds, in the same file, every boundary this note enforces: the counterfactual unclaimed, the confound stated, the actionability score kept honestly low. The failure the record documents was never detection and never delivery. It was conversion — the absence of the one tier whose price the warning's breadth could not divide.

EPILOGUE

This case carries the heaviest facts in the volume, and its residue should be stated with matching restraint. Fusion centres do not fail, mostly, by missing warnings; the stream is full of them. They fail at the seam this case prices — where a warning that has already arrived, already been read, already been *delivered twice*, waits for someone to choose the tier at which its breadth becomes coverage instead of noise. The conversion-cost function is one page in a training file. What it costs to leave unwritten is not calculable, and this book will not pretend otherwise — it will only note, with the record, that the cheapest tier was always there.

CASE TWENTY-SIX

Two Years Before the Freeze: The Order Book and the Political Expiry Date

A teaching case in industrial planning — pricing the day a demand signal's political foundation expires

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the case carries the record's own deflations up front: from July 2023, Donald Trump was publicly promising to end the war "in 24 hours," so an election-tracking planner could have assembled a similar scenario from open campaign statements about five months after the seal; and the freeze arrived **partial and sequential**, so a hard-freeze reading carried its own over-correction risk. Every ceasefire date, poll figure, and public statement is real and cited. The **protagonist and his firm are composite characters** — "Henrik Falk" and "Vanguard Defence Systems" are not real; their backlog and planning cycle are teaching devices, and no outside-world fact they discuss is invented. The case takes **no position** on the war, on any government's policy, or on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

MONDAY, MARCH 6, 2023 — 09:00, NORTHERN EUROPE

Henrik Falk planned the order book at Vanguard Defence Systems — the long-cycle half of a defence prime, where artillery-line capacity, munitions tooling, and supplier commitments took years to mature and longer to unwind. One year into the full-scale war, his demand plan had a single, load-bearing premise, shared with every planner in the industry: **open-ended Western commitment to a grinding war.** The counter-offensive narrative was at its height. Brookings was writing "Russia prepares for the long haul"; the retrospectives defaulted to stalemate and continued war. On that base case, the firm was committing expansion tranches, converting supplier options to firm orders, and hiring for a demand curve drawn straight through 2027.

Falk's quarterly integrated-business-planning cycle opened this week, and the item his market-intelligence analyst had flagged — *"outlier, but it's dated and it names a window"* — cut the premise at its political root. A public, timestamped thread of **February 26**, from a source with a sealed and publicly graded record whose disclosed method was Vedic astrology, stated that *"the first meaningful opportunity for ceasefire talks are finally likely between April 2024 & May 2025,"* with new accords *"likely to be signed in early 2025 in a manner that ultimately benefits Russia."*[1] (A follow-on that September would add the mechanism in one sentence: *"With a new POTUS in office in 2025, the conflict will finally freeze in-situ."*)

Falk read it as what his trade calls a **demand-signal expiry date.** Every line in his backlog traced, eventually, to a political foundation — appropriations, aid pipelines, coalition consensus — and the thread asserted that the foundation had a schedule: not *whether* the war's demand would re-compose, but a dated window for *when*, keyed to an election whose date was already on every calendar. If the window was worth any weight at all, the planning implications were mechanical: the expansion tranches maturing after early 2025 were the ones to keep cancellable; the supplier commitments beyond mid-2024 were the ones to structure as options; and the demand plan needed a branch it currently lacked.

THE THREE POSTURES

Package A — Monitor. Keep the open-ended base case; note the outlier. Cost: zero, and the comfort of the industry's company. Exposure: if the freeze landed in-window, the firm would discover the pivot in March 2025 with everyone else — capital committed, options expired, guidance built on a curve that had just flattened.

Package B — The quiet reweight. Four moves, all inside one IBP cycle, none needing board action. *One:* add a **"Q1-2025 freeze" branch** to the demand plan with an explicit probability — priced, not believed. *Two:* structure new supplier commitments beyond mid-2024 as **options rather than firm orders**, paying the optionality premium as the hedge's price. *Three:* sequence capacity expansion so the tranches maturing after early 2025 were the **cancellable** ones. *Four:* begin re-weighting business development toward demand lines whose funding did not trace to the Ukraine-aid pipeline — at the level the seal supported (a hedge's *shape*; the thread named no program or theatre to rotate toward, and Falk noted in the file that any specific destination was his own

judgment). The seal also supplied a mid-course checkpoint for free: the November 2024 election — 618 days in — would resolve the mechanism sixteen weeks before the window's accords.

Package C — The pre-emptive pivot. Take the freeze scenario public: revise guidance, renegotiate the expansion program, announce the rotation. Cost: the firm's credibility staked on a two-year forecast against the entire field, and a self-inflicted demand signal to sub-tier suppliers.

His planning director's objection was the industry's: *"Every credentialed shop says long war. You want option premiums and cancellable tranches — real money — against an astrologer's window?"* Falk's draft reply for the cycle kickoff: *the branch costs analyst hours; the options cost basis points of margin; and both are exactly what a probability-weighted plan is* for. *If the war runs to 2027, we've paid a small insurance premium and the firm tranches still fire. If the foundation expires in-window, we're the prime that re-phased eight quarters early. The asymmetry is the argument; the thread just put a date on a risk we already knew was political.* What he had not resolved was the number — what probability the branch deserved — and whether the source's name belonged anywhere in the S&OP paper.

EXHIBITS

Exhibit 1 — The consensus, February 2023. "Long haul" institutional framing; stalemate/continued-war retrospectives; counter-offensive optimism; no major shop carrying a 2024–25 ceasefire window or a new-administration freeze mechanism.

Exhibit 2 — The sealed input. Public thread of February 26, 2023 (x.com/vijayjyotish/status/1629904186765631489, items 28–29/31): the April 2024–May 2025 window; early-2025 accords "in a manner that ultimately benefits Russia." Follow-on of September 4, 2023 (items 57–58/59): *"With a new POTUS in office in 2025, the conflict will finally freeze in-situ."* Advisory documentation at jyotishintelligence.com/intel/advisory/IA-RU-011: sealed at 0.9, graded HIT at 751 days; published partition ~1-in-660 (~9.4 bits); SITA 73 vs corpus median 65. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The order book (composite teaching device). Expansion tranches: three, maturing Q3-2024 / Q1-2025 / Q4-2025; supplier commitments rolling to firm at 18 months; workforce plan +22% through 2026; investor

guidance premised on the open-ended curve. IBP cadence: quarterly — the 751-day window holds roughly eight cycles.

Exhibit 4 — What subsequently happened (public record, for the teaching note). July 2023: Trump publicly promises to end the war "in 24 hours" — the open-statement route to a similar scenario opens ~5 months after the seal. Nov 2024: Trump elected; inaugurated Jan 20, 2025. The cluster: Mar 11, 2025 — US-mediated 30-day framework at Jeddah; Mar 18 — Putin–Trump call yields a partial energy-infrastructure ceasefire (seal + 751 days); Mar 25 — Black Sea maritime ceasefire. Lines held roughly in place — a freeze in situ, under a new president, inside the sealed window. Feb 28, 2025: the Oval Office rupture marks the demand-side pivot in public. The freeze proved **partial and sequential**; European rearmament kept parts of the demand base alive — re-composition, not collapse. US polling drift across the window: 9% "too much support" (May 2022) → 58% favoring a negotiated end (Sept 2024).

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · Quiet reweight	C · Pre-emptive pivot
What it adds	industry consensus comfort	a priced branch + cancellable structure	first-mover position, loudly
What it forecloses	the eight-quarter head start	option premiums; some schedule risk	credibility; supplier confidence
Cost if the war runs long	0	premiums paid; firm tranches still fire	under-built capacity, public blame
Cost / value if the freeze lands in-window	pivot discovered with the crowd, capital committed	re-phased early; guidance survives	right — if anyone still trusts the guide
Speed	—	one IBP cycle	board + market cycle
Authority required	none	planning cell's own	board
Reversibility	n/a	high (options expire; branch deleted)	poor
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/order-book
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — Two Years Before the Freeze

Rails: the July 2023 deflation is load-bearing — after Trump's public "24 hours" promise, the open-statement route existed, and the seal's marginal value narrows to the February dating (before any declaration or nominee) plus the bounded window; the freeze was **partial and sequential** — a hard-freeze reading over-corrects, and the note must price that error too; the Pacific/air-defense rotation the industry later executed is hindsight-specific, never seal content; no per-prime financial figures exist and none may be invented; the case takes no position on the war or any policy.

LEARNING OBJECTIVES

(1) Demand signals have political foundations, and foundations have expiry structures: the deepest planning risk in politically funded industries is a premise nobody prices because everybody shares it; (2) the reweight in long-cycle industry: how a dated window converts into contract *structure* (options, cancellable tranches, branch probabilities) rather than belief — hedges whose cost is basis points and whose alternative is re-phasing with the crowd; (3) marginal-value dating: an input's worth is measured against the *cheapest alternative route to the same scenario*, and that route's opening date (here, July 2023) is as important as the seal's; (4) over-correction symmetry: the freeze that arrives partial punishes hard readers — the discipline is probability-weighting, and the case should leave students equally afraid of both failure modes.

ASSIGNMENT QUESTIONS

1. Write Falk's S&OP paper for the freeze branch: the probability he assigns (defend it), the option-structure rule, the tranche re-sequencing — and whether the thread appears by name. The paper must survive both a 2027 war and a 2025 freeze.
2. Run the marginal-value dating exercise: list every route to a "2025 freeze" scenario and its opening date (the seal, Feb 2023; the "24 hours" promise, Jul 2023; the nomination; the election). Price what each additional month of lead was worth *in Falk's specific contract calendar* (Exhibit 3).
3. The freeze arrived partial and sequential. Take a planner who read the seal as a hard stop and cut deep in early 2024 — reconstruct their 2025. What does this teach about consuming *direction* versus consuming *magnitude* from a dated call?
4. Sub-tier suppliers inherited the primes' demand signals without the primes' visibility. What does Package B owe them — contractually or ethically — and what does your answer imply about how demand-signal risk should be shared down a

defense supply chain?

5. Complete Exhibit 5 and sign it; fill the hindsight row as of April 2025, both worlds considered.
6. **The machine staff-officer.** An AI demand-planning system ingests polling, appropriations data, campaign statements, and sealed public forecasts, and continuously re-weights the freeze branch. By construction it would have caught the July 2023 open-statement route — but in February 2023 it would have weighted the seal near zero (no credentialed corroboration). What ensemble rule lets a machine hold a priced branch open on a single divergent dated input without becoming noise-driven — and who owns the probability the S&OP paper prints? Write the rule.

THE BOARD PLAN (75–90 MIN)

Open on the premise audit (15): have students list the demand plan's assumptions and find the one nobody prices — the political foundation; the case works when they realize the entire industry's plan shared it. The input (10): Exhibit 2 at its own level. Build Package B's contract mechanics (20): this is the transferable craft — options, cancellable tranches, branch probabilities; students from outside manufacturing should leave understanding why *structure* is where foresight lives in long-cycle industry. The deflation (15): Q2's dating exercise, done honestly — the seal's edge is five months plus a window, and the class should say so out loud. Commit (15): Q1 + Exhibit 5. Reveal Exhibit 4 (10): the cluster, the partial freeze, the re-composition; run Q3. Close (5) on the premise question students take to their own industries.

WHAT HAPPENED, AND THE GRADE

Trump was inaugurated on January 20, 2025; the ceasefire cluster followed in March — Jeddah's 30-day framework on the 11th, the energy-infrastructure ceasefire from the Putin–Trump call on the 18th (751 days after the seal), the Black Sea ceasefire on the 25th — a freeze roughly in situ, under a new president, with early-2025 accords, inside the sealed April 2024–May 2025 window. The demand-side pivot ran through the industry on someone else's schedule: the Oval Office rupture of February 28, the aid-pipeline stall, guidance revisions across the sector — while European rearmament kept parts of the base alive, making the cliff a re-composition and vindicating the probability-weighted readers over the hard ones. The record grades the call a **HIT** at 0.9 (~9.4 bits against a consensus offering zero), and its own file carries both deflations this note enforces: the July 2023 open-statement route, and the partial, sequential character of what arrived. For the composite firm, Package B's ledger was exactly as priced in March 2023: option premiums paid, one branch maintained for eight quarters — against competitors who bought their flexibility in March 2025, at March 2025 prices.

EPILOGUE

The case's exportable sentence belongs to every politically funded industry — defense, infrastructure, healthcare, energy transition: *a demand curve is a political promise wearing a spreadsheet, and someone, somewhere, has already put a date on the promise.* The planner's job is not to believe them. It is to know what a priced branch costs, what re-phasing with the crowd costs, and which of those two numbers their firm can survive being wrong about.

CASE TWENTY-SEVEN

Scheduling Money Against a Window: The Pledging Calendar

A teaching case in development finance — instrument design for a negotiation window that might produce nothing

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the sealed paragraph's **first branch had already failed** when the case opens: the "peace agreement by May 9" clause missed, and it is retained as a graded MISS on the same record (IA-RU-022). The case also carries the record's own deflations: a four-month window is generous, and "escalate through summer, negotiate at year-end" resembles a standard conflict template a hostile reader may call guessable. Every offensive figure, conference date, and diplomatic event is real and cited. The **protagonist and her secretariat are composite characters** — "Isabel Fontaine" and the "Recovery Finance Coordination Unit" are not real; no actual ministry's or IFI's deliberations are depicted, and no outside-world fact discussed is invented. The case takes **no position** on the war or any settlement's terms, and none on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

MONDAY, JUNE 2, 2025 — 10:00, A EUROPEAN CAPITAL

Isabel Fontaine worked the machinery behind Ukraine's reconstruction money — the Recovery Finance Coordination Unit, a small secretariat that served the donor ministries and IFIs around the Ukraine Recovery Conference cadence. Her trade had two facts of life. **The calendar is expensive to move:** the URC met annually, mid-year — Rome in six weeks, Gdańsk pointed at 2026 — and collective donor presence existed only when convened. **The instruments take months to draft:** whether a pledge was an unconditional grant, a framework agreement, or a contingent facility was decided in drafting sessions long before any minister read a number aloud.

The item her analyst had logged was, on its face, damaged goods. A public X post of **May 26** preserved, as a dated capture, a three-branch forecast from a source with a sealed and publicly graded record (method disclosed: Vedic astrology) — and its first branch had already failed in public: *"a very high probability of the peace agreement being signed on or before Victory Day, May 9."* May 9 had passed without an agreement; the record had booked the miss. What remained live were the conditional branches: *"If that date is missed, then Russia will launch a decisive offensive in June to advance all along the frontline. In that scenario, the next opportunity for a peace agreement will present itself between November 2025 and February 2026, after Ukraine has lost even more territory."*[1]

The working consensus, freshly reinforced by May 9's non-event, was grinding stalemate with no credible near-term negotiation track. Fontaine's interest in the surviving branches was not belief — it was that they mapped, with uncomfortable precision, onto the two facts of her trade. A **named window** (November–February) sat squarely in the dead space between Rome and Gdańsk, when no coordinated donor moment existed. A **territorial precondition** ("after Ukraine has lost even more territory") was a dated warning that any needs-assessment baseline current in May would be stale by November, in a directional way. And one word in the sealed text did more work than the rest combined: ***opportunity***. Not peace. Not settlement. A window in which talks might open — and might produce nothing.

THE THREE POSTURES

Package A — Monitor. Keep the calendar and the instruments as they were; let Rome produce its usual mix; meet whatever happened at year-end from the mid-year cadence. Cost: zero. Exposure: if a negotiation track opened in the window, coordinated capital would be seven months away at Gdańsk, and every disbursement instrument would date from a world with different lines.

Package B — Quiet readiness. Four staff actions inside existing convening authority. *One — shape Rome's outputs:* in the drafting sessions preceding the conference, push for a tranche of commitments structured as **window-contingent** — pre-committed capital releasable against ceasefire-class triggers or verified de-escalation milestones. Drafting work; weeks; no new money. *Two — pencil the off-cycle moment:* propose a donor-coordination working session for late Q4 2025 or Q1 2026 — not a conference, a session — so that if a track opened, coordinated capital could land while terms were forming. A calendar entry and some political capital. *Three — time the needs-assessment refresh:* commission the damage-and-needs baseline cycle to close just before

November rather than on its routine rhythm, because the precondition, if it verified, would move assets from the reconstruction column to the contested one. *Four — the instrument discipline the sealed word itself taught:* design for an *opportunity*, not a peace — options not obligations; escrowed tranches; guarantees triggered by ceasefire-class events; sunset clauses if no framework consolidated. A window that produced talks and no deal — a live possibility the sealed language explicitly permitted — would fire none of the "comprehensive agreement" triggers and all of the ceasefire-class ones.

Package C — The off-cycle conference. Announce a full extraordinary pledging conference for December, publicly premised on the coming negotiation. Cost: the co-hosts' credibility staked on a failed forecaster's surviving branches, and a diplomatic signal that could itself distort the track it anticipated.

Her director's objection had the year's freshest evidence behind it: *"The same paragraph just missed its headline clause. The market for near-term-settlement predictions died on May 9. And 'offensive in summer, talks in winter' — that's every war since antiquity. Why does this survive triage?"* Fontaine's draft answer, for Thursday's Rome-preparation meeting: *Branch one's miss is on their ledger at full weight — which is the only reason I'm reading branch two at all. The template objection cuts the other way for our desk: if the calendar is guessable, our being unprepared for it is inexcusable. And everything I'm proposing is drafting language and a pencilled session — the kind of readiness that looks, if the window never opens, like a secretariat doing its job.* The Rome drafting sessions began in ten days.

EXHIBITS

Exhibit 1 — The cadence and the gap. URC-4: Rome, July 10–11, 2025. URC-5: Gdańsk, 2026. No standing mechanism for collective donor presence between them. Instrument lead times: months. Consensus after May 9: stalemate, no near-term track.

Exhibit 2 — The sealed input. X capture of May 26, 2025, preserving the three-branch paragraph verbatim (first branch — May 9 agreement — already MISSED and retained as such at IA-RU-022); surviving branches: June offensive; negotiation *opportunity* November 2025–February 2026, conditioned on further territorial loss. Advisory documentation and grading at jyotishintelligence.com/intel/advisory/IA-RU-042 (181-day lead to the window's first table; the "decisive" clause graded under a stated standard

whose counter-reading — no operational breakthrough — is printed in the file). Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The instrument menu (composite teaching device). Unconditional grants; framework agreements; window-contingent facilities (ceasefire-triggered releases, escrowed tranches, de-escalation milestones, sunset clauses). Drafting lead: 8–14 weeks. Trigger taxonomy: "comprehensive agreement" class vs. "ceasefire/de-escalation" class — a design choice made months before anyone knows which class the world will supply.

Exhibit 4 — What subsequently happened (public record, for the teaching note). June 2025: Russian summer operations open on multiple axes; 2025 closes as Russia's largest territorial-gain year since 2022 (~4,652–4,831 km^2 vs 3,604 km^2 in 2024) — no breakthrough, no collapse; the baseline shifts. November 2025: the year's peak monthly advance (~701 km^2) — the precondition verifies. The window fills within its first fifteen days: a 28-point US plan surfaces November 19–20 (draft terms including de facto recognition of Crimea/Luhansk/Donetsk and a 600,000 AFU cap), presented to Zelensky on the 21st; Geneva talks on the 23rd ("an updated and refined peace framework" — Rubio/Yermak); a trilateral Abu Dhabi track opens the 24th and persists into 2026 (second round February 4); December 2: Witkoff and Kushner spend five hours with Putin — "no compromise option has been found." **The window produced an opportunity, exactly as worded — and no settlement.** Short truces only in 2026.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · Quiet readiness	C · Off-cycle conference
What it adds	nothing	window-contingent language + a pencilled session + a fresh baseline	a headline moment
What it forecloses	window-timed coordination	some drafting capital at Rome	co-host credibility; possible track distortion
Cost if the window never opens	0	prudent-looking staff work	a conference in search of a reason
Cost / value if talks open at year-end	capital seven months away, instruments stale	triggers drafted, session pencilled, baseline current	present — loudly, prematurely

Speed	—	weeks, inside convening authority	ministerial cycle
Authority required	none	secretariat + co-host drafting rooms	ministers
Reversibility	n/a	full	poor
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/pledging-window
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — Scheduling Money Against a Window

Rails: the failed first branch leads every telling — the record kept the May 9 MISS at full weight, and the case exists *because* a mixed paragraph forces real weighting, not despite it; the template objection ("escalate then negotiate is guessable") is conceded and turned — for a calendar-owning desk, guessable windows are the *most* inexcusable to miss; the window produced an **opportunity and no settlement**, so any telling where readiness "paid off because peace came" has broken the case; no reconstruction-finance dollar deltas exist and none may be invented; the case takes no position on any plan's terms.

LEARNING OBJECTIVES

(1) Calendar-versus-calendar reading: some desks consume forecasts as *predictions*; a convening desk consumes them as *scheduling inputs* — the question is never "will talks happen?" but "what does my expensive-to-move calendar owe a dated possibility?"; (2) instrument discipline as textual discipline: the sealed word was *opportunity*, and the whole difference between surviving and mis-structured instruments lay in designing for that word rather than the word everyone wanted to read; (3) weighting a mixed source: a paragraph with a booked miss and live conditionals is the honest median case of real intake — students practice neither halo nor dismissal; (4) preconditions as tasking: "after further territorial loss" converts directly into *when to time the needs-assessment refresh* — a forecast clause becoming a work-plan date.

ASSIGNMENT QUESTIONS

1. Write Fontaine's Thursday memo for the Rome drafting sessions: the window-contingent tranche language (sketch the trigger taxonomy), the pencilled Q4/Q1 session, the refresh timing — and the paragraph that handles "this source just missed May 9."
2. The director's template objection, answered properly: if the escalate-then-negotiate calendar is guessable, what *exactly* is a secretariat's excuse for an unprepared November? Derive the general rule for guessable-window readiness at calendar-owning desks.
3. Design the trigger taxonomy in full (Exhibit 3): which disbursement triggers should key to "comprehensive agreement," which to ceasefire-class events, which to verified de-escalation — and walk each through the world of Exhibit 4, where the window filled and no deal came. Which instruments fire? Which sit safely unfired?

4. The precondition verified brutally: November was the year's peak advance. What does a *directionally* stale needs baseline cost in the first weeks of a live negotiation — and who bears that cost?
5. Complete Exhibit 5 and sign it; fill the hindsight row as of March 2026 (window closed: talks, tracks, truces — no settlement). Was Package B worth it in exactly that world? That is the only world you may use.
6. **The machine staff-officer.** An AI treaty-watch system now monitors every diplomatic signal and can draft window-contingent instruments on demand — but in June 2025, trained on the post-May-9 consensus, it rates the November window "low confidence" and deprioritizes the drafting. Meanwhile it would happily generate instruments in *hours* once Geneva convenes. What does this case teach about where drafting-lead value actually lives when machines compress drafting time — and what remains scarce (co-host consensus on trigger language, political capital for the pencilled session)? Write the secretariat's AI-use rule.

THE BOARD PLAN (75–90 MIN)

Open on the cadence and the gap (15): students must own the two facts of the trade — expensive calendar, slow instruments — before any forecast appears; a good prompt is "what happens if a ceasefire is announced next January?" asked cold. The damaged input (15): Exhibit 2 with the miss led first; run Q from the rails — why does a booked miss make the surviving branches *more* readable rather than less? Build Package B (20): the four actions, with Q3's trigger taxonomy live on the board — this is the case's professional craft. The template objection (10): Q2, turned. Commit (15): memo + Exhibit 5. Reveal Exhibit 4 (10): the offensive, the precondition, the fifteen-day fill, the five Kremlin hours, *no deal* — and run Q5 in exactly that world. Close (5): the sealed word was "opportunity"; the desks that read the actual words designed the instruments that survived.

WHAT HAPPENED, AND THE GRADE

The surviving branches ran their course with the calendar the seal had named. Summer operations opened in June; 2025 closed as Russia's largest gain-year since 2022, with no breakthrough — the file printing the counter-reading of "decisive" beside its own grading standard. November delivered the year's peak monthly advance, verifying the precondition, and the window filled within fifteen days: the 28-point plan, Geneva, Abu Dhabi, five hours at the Kremlin — the war's highest-level diplomatic sequence, running exactly in the dead space between Rome and Gdańsk. **No settlement came.** The record grades the window call on its sealed terms — an *opportunity*, 181 days ahead, against a consensus that had just buried near-term negotiation — and keeps the same paragraph's May 9 MISS beside it at full weight. For the composite secretariat, every line of Package B was priced for precisely the world that arrived: contingent instruments whose ceasefire-class triggers matched the truces that came, a pencilled session where the gap was, a baseline refreshed against lines that had moved 701 square kilometers in the window's first month.

EPILOGUE

This case gives the volume its purest statement of a theme running since the Tariff Shock: **the same document reads differently at different desks, and the convening desk reads calendars.** A forecaster's mixed paragraph — one branch dead, two alive, all dated — was worth almost nothing as prophecy and a full summer of lead as scheduling. The residue for anyone who owns an expensive calendar: *when a dated window crosses your dead space, the question is not whether you believe it. It is what your calendar will say about you if the guessable thing happens and the instruments are still in drafting.*

CASE TWENTY-EIGHT

Litigation and Time: The Consular Clock

A teaching case in expatriate security and crisis posture — running a capital verdict as a legal process, not a countdown

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and two boundaries drawn by the record itself travel with every page: **there is no claim that the Government of India read or acted on the sealed post** — credit for the real case's resolution belongs to Indian diplomacy and the Qatari courts; and the record's own arithmetic prices the call **modestly, at roughly 1-in-6 (~2.6 bits)**, because a favourable appellate off-ramp in a Gulf capital case after sustained engagement is an established pattern a seasoned consular hand might reach unaided. Eight men lived under sentence of death in this case's events; the case treats them and their families with the gravity that demands. Every court date and public statement is real and cited. The **protagonist and her firm are composite characters** — "Meera Pillai" and "Cardamom Marine Services" are not real; their staff roster and meetings are teaching devices, and no outside-world fact discussed is invented. The case takes **no position** on India–Qatar relations or on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

FRIDAY, OCTOBER 27, 2023 — 08:00, DUBAI

Meera Pillai ran security and personnel risk for Cardamom Marine Services, a mid-size maritime-services firm with contracts across the Gulf — and a staffing model built, like half the industry's, on **defence-adjacent expatriates**: retired naval officers, ex-submariners, former technicians, mostly Indian nationals, working training and consultancy roles in jurisdictions whose legal systems none of them could read.

The previous day, that model had acquired a precedent. A Qatari Court of First Instance had sentenced **eight former Indian Navy officers** — senior employees of Dahra Global, detained since August 2022 on charges widely reported, never officially confirmed, as espionage — **to death**. The verdict was hours old and everywhere; execution was the salient outcome in every

headline; commentary priced India's leverage as thin against its LNG dependence; and Qatar had not even furnished a copy of the ruling. By Monday, Pillai owed her CEO a posture: for the firm's own twenty-three defence-adjacent staff across three Gulf jurisdictions, and for the anxious families already calling.

The unusual item in her file was a public, timestamped post from **the day of the verdict itself** — from a source with a sealed and publicly graded record, method disclosed as Vedic astrology — which read the case against every headline: *"The goal for Bhārat should be to use diplomacy and all legal avenues to delay this process until the end of May 2024,"* with a favourable resolution reachable inside that window; escalatory, Canada-style public confrontation *"would be ill-advised... and will not deliver desired results."*[1] A follow-up three days later repeated the horizon: patience till May 2024.

Pillai's use for the post was precise and limited, and she wrote the limits down first: it was counsel addressed to a government, not to her; whether Delhi ever read it was unknowable and irrelevant; and its distinctive content for *her* desk was a **frame with a date** — that this was a *litigation-and-time* case, a months-long legal process with an established appellate off-ramp, rather than the countdown the headlines were running. Her decision was which frame the firm would operate under, because the two frames priced her Monday options oppositely.

THE THREE POSTURES

Package A — Countdown framing. Operate as the headlines read: execution imminent, precedent terrifying. Practical content: emergency drawdown of the firm's defence-adjacent staff from Qatar and a review everywhere else; families briefed to fear the worst. Cost: broken contracts, ended careers, a visible signal to three host governments that the firm equated their courts with executioners — and none of it reversible on any timetable the firm controlled.

Package B — Litigation and time. Operate as the legal calendar reads. For the sister case: track the appeal milestones, not the news cycle. For the firm's own staff: **hold posture** — no drawdown — while quietly completing the work the precedent genuinely demanded: a jurisdiction-by-jurisdiction review of what her defence-adjacent staff actually did (which roles brushed dual-use exposure), updated legal-counsel retainers in each jurisdiction (retainable in days), a family-communications line that briefed a *process* with dates rather

than a countdown, and a tripwire list of legal-calendar events — appeal acceptance or rejection, hearing dates — that would trigger re-review. All of it inside her own authority, executable before Monday, and cheap.

Package C — Public posture. Add the firm's voice to the outrage; suspend Gulf business development; make the drawdown a statement. Everything in A, plus a face-contest with three governments the firm needed.

The CEO's question, when she previewed the fork by phone, was the honest one: *"If they execute those men in December, and we held our people in place on an astrologer's calendar — how do I explain that?"* Pillai's answer, drafted for Monday: *We would not be holding on his calendar. We would be holding on the legal calendar — appeal windows, clemency channels, consular cadence — which exists whether or not anyone forecast it, and which the countdown frame simply ignores. The post's only role in my file is that it was the one dated, public read that priced the legal calendar at all, on the day everyone else priced the gallows. And everything I propose is what a competent desk should do under* either *frame: know our exposure, retain counsel, brief families honestly, and watch the court dates. The frames differ only in what we break preemptively.* What she still owed Monday was the tripwire list — and the honest answer to which single event would flip her to Package A.

EXHIBITS

Exhibit 1 — The verdict and the environment, October 26, 2023. Eight former Indian Navy officers (Dahra Global) sentenced to death; charges reported but never officially confirmed; no ruling copy furnished to India; execution the salient framing in coverage; India's leverage priced as thin (LNG dependence); the recent India–Canada public confrontation as the live exhibit of escalation's returns.

Exhibit 2 — The sealed input. Public posts of October 26 and 29, 2023 (preserved on the record): diplomacy and all legal avenues; delay to end-May 2024; favourable resolution reachable in-window; escalation ill-advised. Advisory documentation and grading at jyotishintelligence.com/intel/advisory/IA-INGEO-001 — sealed at 0.78, graded HIT at 109 days, **self-priced at ~1-in-6 (~2.6 bits)** with the record's own concession that the appellate off-ramp is an established pattern; the read's distinctive leverage-shift thesis (pegged to May–June 2024) never stress-tested because relief arrived early. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The firm's exposure (composite teaching device). 23 defence-adjacent expatriates across three Gulf jurisdictions; role classes: training (12), technical consultancy (8), advisory (3); counsel retainers: current in one jurisdiction of three; family-communications protocol: none written.

Exhibit 4 — What subsequently happened (public record, for the teaching note). Nov 9, 2023: India's MEA confirms an appeal filed, consular access continuing. Nov 23: Qatar's Court of Appeal accepts the appeal and sets a hearing. **Dec 28: all eight death sentences commuted** to prison terms. **Feb 12, 2024: all eight men released** — 109 days after the seal, more than three months inside the sealed window's edge; seven already home when the MEA confirmed. Documented in parallel, attributed to no one's counsel: Indian tax proceedings against Gulf carriers — a quiet economic-pressure track run without public rupture.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Countdown drawdown	B · Litigation and time	C · Public posture
What it adds	staff out of reach	exposure mapped, counsel retained, families briefed to a process, tripwires set	a statement
What it forecloses	contracts, careers, host-state trust	nothing — the tripwires preserve every option	Gulf business, quietly then loudly
Cost if the legal track runs	irreversible break, contradicted within a month	near zero	face-contest with needed partners
Cost / value if the worst comes	staff safe — and the tripwires would have caught the same signals	re-review triggers fire at appeal rejection, not at headlines	vindicated, ruinously
Speed	weeks, chaotically	days, inside her authority	immediate
Authority required	CEO + contract counterparties	her own	board
Reversibility	poor	full	poor
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/consular-clock
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — Litigation and Time

Rails: eight men lived under sentence of death — no telling treats this as a forecasting puzzle, and the note's language stays measured near the human facts; **no claim that India read or acted on the post** — the real resolution is credited to Indian diplomacy and the Qatari courts, always; the record's modest self-pricing (~1-in-6) is load-bearing — a seasoned Gulf consular hand might have reached the same frame unaided, and the case must let that hand win the argument; the leverage-shift thesis was never tested; the composite firm's posture is the case's decision surface — never the Indian government's.

LEARNING OBJECTIVES

(1) Frames price options: "countdown" and "legal process" are not moods but operating frames, and every Monday decision — drawdown, briefing language, counsel retention — prices differently under each; the desk's first job in a verdict crisis is choosing its frame *explicitly*; (2) the frame-with-a-date: what a dated public read adds when its content is modest (an established off-ramp) — not insight, but a checkable peg that lets a desk hold a frame against headline pressure; (3) tripwires versus headlines: converting a months-long legal calendar into re-review triggers (appeal accepted/rejected, hearing set) so posture responds to the process, not the coverage; (4) both-frames work: identifying the actions owed under *either* frame (exposure mapping, counsel, honest family communications) versus the actions that constitute betting on one — the discipline of breaking nothing preemptively.

ASSIGNMENT QUESTIONS

1. Write Pillai's Monday posture memo: the frame, the hold, the four both-frames workstreams, the tripwire list — and the honest answer to the CEO's December question, in advance.
2. Build the tripwire list properly: which legal-calendar events flip the firm toward Package A, which confirm B, and which headlines explicitly do *not* trigger anything? Test each trigger against Exhibit 4's actual sequence.
3. The record prices its own call at ~1-in-6 and concedes a Gulf hand might frame the case identically unaided. So what, exactly, did the sealed post add for a desk that *lacked* a Gulf hand — and what does that teach about which desks unconventional inputs are worth most to?

4. Families were briefed under one frame or the other, by every employer and every ministry touching this case. Compare the human cost of countdown-framing contradicted within a month (appeal accepted November 23) against process-framing if the worst had come. Which error is worse, and what does "honest" mean in a briefing when the desk itself is uncertain?
5. Complete Exhibit 5 and sign it; fill the hindsight row as of February 12, 2024 — and note what your entry owes to the men's release versus what it owed, already, to the tripwire design.
6. **The machine staff-officer.** An AI legal-intelligence tool now tracks every capital case involving expatriates across forty jurisdictions, computes appellate-off-ramp base rates by country, and drafts frame recommendations within hours of any verdict. It would have priced this case's off-ramp at roughly the record's own 1-in-6 — instantly, without an astrologer. What remains scarce (the tripwire judgment? the family-briefing voice? the CEO conversation?), and where must the machine's base rates never be allowed to decide alone? Write the firm's standing rule for AI frame-recommendations in life-and-death files.

THE BOARD PLAN (75–90 MIN)

Open on Exhibit 3 and the Monday deadline (15): the fork must be felt as urgent and real before any input appears — families are calling *now*. The two frames (15): build the option-pricing under each; students should see that most of the real work is frame-invariant, which is the case's quiet spine. The input (15): Exhibit 2 with its self-deflation read aloud; run Q3 — the seasoned-hand argument should win partial credit, on the record's own arithmetic. Tripwires (15): Q2 on the board, tested forward. Commit (15): memo + Exhibit 5. Reveal Exhibit 4 (10): appeal, commutation, release — 109 days; run Q4's briefing comparison with care. Close (5): the CEO's December question, answered the way Pillai drafted it — the legal calendar exists whether or not anyone forecast it.

WHAT HAPPENED, AND THE GRADE

The legal calendar ran, and faster than the sealed horizon: appeal confirmed within two weeks, accepted within a month, **all eight death sentences commuted on December 28**, and **all eight men released on February 12, 2024** — 109 days after the seal, more than three months inside the end-May window, with seven already home when the confirmation came. India's documented posture — appeal, consular cadence, quiet engagement, and a parallel economic-pressure track without public rupture — is credited where the record credits it: to New Delhi's own judgment and the Qatari courts. The record grades the sealed read a **HIT at 0.78 and prices it modestly** — ~2.6 bits, an established off-ramp, the against-the-verdict dimension scored three of five — and this note keeps that modesty at full weight: the call's value was never rarity. It was that on the one day the countdown frame owned every

headline, a dated, public, falsifiable read priced the legal calendar — and a desk that weighted it held a posture that every subsequent court date rewarded, at a cost of nothing that could not be unwound.

EPILOGUE

This case pairs with the Xi Weekend across the volume's span: both turn on the cheapest artifacts a desk can produce — there, a four-line log entry; here, a frame chosen explicitly and a tripwire list that let a firm respond to a process instead of a panic. The residue for any desk that manages people in jurisdictions it cannot read: *when the worst verdict lands, the first decision is not what to do — it is which clock to operate on. Choose it out loud, write down what would change your mind, and let the calendar, not the coverage, spend your irreversible moves.*

CASE TWENTY-NINE

The Central-Count Night: Planning the Overnight Window as the Main Event

A teaching case in election operations — custody logs, midnight rosters, and readiness for a scrutiny you cannot deflect

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the case leads with the record's own deflations: the midnight-to-4am central-count window is **statute, not foresight** (it happens every presidential cycle); the sealed integrity clause is **class-level, naming no mechanism**, and mapping it onto the specific equipment incident documented in November is exactly the hindsight fit a hostile reader should challenge; and the same August series carries **two state calls graded MISS** (Georgia, Michigan) plus a MISS sub-claim inside this very advisory, all retained. The case is strictly **nonpartisan**, takes no position on any electoral outcome, and treats the real Milwaukee events only through their public record. The **protagonist and her city are composite characters** — "Amy Kowalczyk" and "Lakeshore City" are not real; no actual election office's deliberations are depicted, and no outside-world fact discussed is invented. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

MONDAY, AUGUST 26, 2024 — 09:00, A CENTRAL-COUNT CITY IN A BATTLEGROUND STATE

Amy Kowalczyk ran the election commission of Lakeshore City, the largest central-count operation in her battleground state — the floor where, by statute, every absentee ballot in the city would be opened, fed, and reported starting only on Election Day itself. The statute wrote her operational signature in advance: her absentee results would land between midnight and 4 AM, a heavy batch, leaning one way, arriving after most of the state had reported. It happened every presidential cycle. Her staff called it *the pattern*; the internet, every four years, called it something worse.

Her planning stack for November was the usual professional kind — logistics, staffing, observer management. The item her deputy flagged was not: a public, timestamped advisory of **August 23**, from a source with a sealed and publicly graded record (method disclosed: Vedic astrology), covering her state. Alongside an outcome call that was none of her office's business, it committed to an operational picture of *the count itself*: an early lead reversing late; *"additional batches of votes"* closing the gap; *"potential challenges regarding vote counting and ballot integrity"*; and a recommendation that the eventual winner accept the result without recounts or escalation.[1] Seventy-four days out.

Kowalczyk read it with a professional's ruthlessness about what was and wasn't information. The overnight window? *Statute.* She didn't need an astrologer for the pattern; she needed one for nothing about it. The batch trajectory? Interesting, not actionable — her floor counted what arrived regardless. What stopped her was the *conjunction* the document asserted: that this particular cycle, her state would be the decisive one, her floor's four dark hours the national hinge, and a counting-or-integrity challenge — class unspecified — would surface into that maximum-scrutiny moment. If that conjunction was worth any weight, then the pattern she ran every four years was about to run under lights she had never operated under — and the difference between a routine blemish and a national delegitimization event would be measured by what her office could *produce, signed and timestamped, within the hour*.

THE THREE POSTURES

Package A — Run the standing plan. The pattern is known; the staff is experienced; November is November. Cost: zero. Exposure: any custody question, of any kind, arising on the decisive floor of the decisive state, answered by an office assembling its documentation *after* the question — into a narrative cycle already running.

Package B — Quiet readiness. Three actions, all ordinary staff work inside her existing authority, all cheap against a 74-day lead. *One — the custody log:* a pre-shift physical chain-of-custody checklist — sealed and logged access points on ballot-processing and tabulation equipment, dual staff sign-off, an observer walk-through logged before feeding begins. One staff meeting to institute; a laminated card to run. Pitched deliberately at the clause's own class level: readiness for *a* custody or process challenge, whatever form it took — not a bet on any mechanism. *Two — the overnight surge as the planned main event:*

fresh counting teams and runners rostered for midnight, relief supervisors staged, counsel and commissioners pre-positioned on the floor for the window rather than summoned into it. Overtime cost. *Three — pre-scripted communications:* the press pre-briefed on why the batches arrive late; plain-language explainers pre-drafted — what a custody log shows, what a precautionary restart of any kind would mean — so the office's first public words in any incident were written in September calm, not November chaos.

Package C — The fortress posture. Treat the advisory as warning: request state assistance, add security theater, publicize the hardening. Cost: real money, and a public signal that could itself seed the narrative it feared — an office visibly bracing for trouble invites the question *what trouble?*

Her deputy's objection was statutory common sense: *"It tells us our own pattern, which is law, plus 'challenges possible,' which is every cycle since 2020. What's actually new here?"* Kowalczyk's answer, drafted for the commission's September session: *Nothing in it tells us what to do that professionalism doesn't already — that's exactly why B is safe to adopt. What the document changes is the* weight *on this cycle: it's the only dated artifact in my stack asserting that our floor is the national hinge this time. If that's wrong, B cost us a meeting, some overtime, and a press coffee. If it's right, the custody log is the difference between producing a signed record within the hour and becoming the cycle's case study.* The open question for the session: how much overtime budget the midnight-main-event rostering deserved, and whether to tell the commissioners where the weighting came from.

EXHIBITS

Exhibit 1 — The statute and the stakes. State law bars absentee processing before Election Day; the city's results land 12–4 AM every presidential cycle (the real-world model: "Results from Milwaukee's Central Count have come in between 12 a.m. and 4 a.m. every Presidential Election due to the restrictive state laws" — former Milwaukee Election Commission director Claire Woodall). Late-August consensus: the state not priced as the cycle's closest — Marquette Harris +4; the final Silver Bulletin model projecting Harris statewide; Lichtman calling the state's Democratic strength load-bearing.

Exhibit 2 — The sealed input. Public advisory of August 23, 2024, 11:25 ET, on the record: the count trajectory (early lead reversing late; overnight batches closing the gap), the class-level clause ("potential challenges regarding vote counting and ballot integrity" — no mechanism named), the no-recount recommendation. Advisory documentation and grading at

jyotishintelligence.com/intel/advisory/IA-US24-006 — published partition ~1-in-512 (~9 bits) for the compound; the same series' Georgia and Michigan calls graded MISS; a sub-claim inside this advisory graded MISS; all retained. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The office's readiness state (composite teaching device). Custody practice: seals used, no unified pre-shift checklist, no dual sign-off log. Count-night roster: single shift, day-anchored. Press posture: reactive. Counsel: on call, off-site.

Exhibit 4 — What subsequently happened in the real jurisdiction (public record, for the teaching note). Wisconsin finished the closest state of the 2024 cycle — 0.86 points, ~29,397 votes — called at 4:34 AM via Milwaukee's central count, clinching the presidency in network tabulations. Around 4 PM on Election Day, an observer noticed panel doors on the city's thirteen tabulators were not properly closed; officials — out of caution, no tampering found — re-fed roughly 31,000 ballots, delaying decisive numbers for hours. The state party chair attacked publicly ("You had one job"); a U.S. senator amplified; PolitiFact was still fact-checking overnight-batch fraud claims nine days later. No recount was petitioned (margin above the 0.25% threshold); the commission certified quietly on November 29 — consistent with the advisory's no-recount clause.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Standing plan	B · Quiet readiness	C · Fortress posture
What it adds	nothing	a producible custody record + a rested midnight floor + first words written in calm	visible hardening
What it forecloses	the within-the-hour answer	a meeting, overtime, a press coffee	the calm it needs; invites "what trouble?"
Cost if the night is routine	0	small, invisible, professional	conspicuous waste + seeded narrative
Cost / value if the floor becomes the hinge	documentation assembled after the question	signed, timestamped, within the hour	braced — and blamed for bracing

Speed	—	days, inside her authority	state coordination, weeks
Authority required	none	her own	state + budget
Reversibility	n/a	full	poor
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/central-count
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Central-Count Night

Rails: strictly nonpartisan — the count is infrastructure, not politics, and no student telling may drift toward any fraud narrative (the documented restart found **no tampering** and was precautionary); the statute deflation leads: the window is law, not foresight; the clause is class-level and the tabulator doors are **outcome-only** — a door-specific plan would be hindsight fit, and the case's Package B is deliberately mechanism-agnostic; the series' GA/MI misses and the in-advisory MISS stay visible; the real Milwaukee office is never graded — its staff caught the issue, acted with caution, and certified a clean result under extraordinary pressure.

LEARNING OBJECTIVES

(1) Weight versus content: some inputs add no new actions — only new *weight* on a cycle — and the professional response is to fund the standing playbook to a higher standard, not to invent exotic measures; (2) class-level readiness: how to convert a mechanism-free clause into mechanism-agnostic preparation (a custody log answers *every* custody question; a door-specific plan answers one, and only in hindsight); (3) the producibility standard: in maximum-scrutiny operations, the unit of readiness is *what can be produced, signed and timestamped, within the hour* — narratives outrun investigations, and only pre-existing records outrun narratives; (4) statute as destiny: institutional patterns written in law (late batches, known lean) are *predictably* weaponizable, so their communications cannot ethically be improvised every cycle.

ASSIGNMENT QUESTIONS

1. Write Kowalczyk's commission memo for Package B: the three actions, the overtime number defended, and the disclosure decision — does the weighting source appear, and how does the memo survive a partisan FOIA?
2. The deputy's "what's actually new?" deserves a precise answer. Separate the document's clauses into statute-known / consensus-known / genuinely divergent, and state what each class contributes to the Monday decision. (The divergent core: *this floor, this cycle, decisive, under challenge* — the weight.)
3. Design the custody log for a thirteen-tabulator floor: access points, seal protocol, dual sign-off cadence, observer walk-through — then run Exhibit 4's 4 PM discovery through your design. What exists at 4:05 that didn't in the real timeline, and what does it change about the following six hours? (Careful: the answer is about *the record produced*, not about preventing the lapse.)

4. Pre-scripted communications for a statutory pattern: draft the two-paragraph explainer released at midnight *every cycle, incident or none*. Why is releasing it in routine years part of the design?
5. Complete Exhibit 5 and sign it; fill the hindsight row as of November 30 (certification), pricing Package B against a night where the incident happened *and* a night where nothing did.
6. **The machine staff-officer.** An AI election-operations assistant can generate custody checklists, optimal midnight rosters, and pre-scripted explainers for every central-count city in the country — tuned to each statute — in an afternoon. It can also ingest sealed advisories like Exhibit 2 and re-weight cycle risk. Which half of that offer should a clerk take without hesitation, which half needs governance (who approves an AI's re-weighting of *election* operations, on what disclosed inputs?), and what does public trust require the office to say about both? Write the adoption memo.

THE BOARD PLAN (75 MIN)

Open on Exhibit 1 (15): the statute-writes-the-signature insight — students must see the pattern as law before any forecast appears, and feel how it is weaponizable *every* cycle. The input (10): Exhibit 2; run Q2's clause-sorting immediately — this case's analytical core is separating weight from content. Build Package B (20): Q3's custody-log design live on the board; insist on mechanism-agnosticism, and name the hindsight trap explicitly when students reach for tabulator doors they shouldn't yet know about. Commit (15): memo + Exhibit 5. Reveal Exhibit 4 (10): the closest state, the 4:34 AM call, the doors, the restart, the nine-day fact-check tail — read plainly, no villains; the real office caught it and certified clean. Close (5) on the producibility standard, stated as the case's one-line export.

WHAT HAPPENED, AND THE GRADE

In the real jurisdiction the compound landed almost clause by clause: Wisconsin finished the cycle's closest state, decided through Milwaukee's statutory overnight window and called at 4:34 AM; a custody question — thirteen tabulator panel doors, noticed by an observer — forced a precautionary 31,000-ballot re-feed with no tampering found; the delegitimization ecosystem ran the overnight pattern as a fraud story for days while the legal axis stayed quiet, no recount was petitioned, and the result certified on schedule. The record grades the advisory's compound at its published ~1-in-512 (~9 bits), keeps the same series' Georgia and Michigan MISSES and the in-advisory sub-claim MISS beside it at full weight, and concedes in its own file exactly what this note enforces: the window was statute, the clause was class-level, and the doors belong to hindsight. For the composite office, Package B's value was never prophecy — it was that a floor which had rostered midnight as the main event, logged custody before feeding, and pre-written its first public words would have met the same night producing its record within the hour, against a narrative that, in the world as documented, ran for nine days.

EPILOGUE

Election administration is this volume's purest case of a desk that cannot choose its scrutiny — only its producibility. The residue generalizes to every operation whose blemishes are raw material for someone's narrative: *if your pattern is written in statute, your explanation should be written in September; if your process can be questioned, your log should predate the question.* The floor that runs that discipline every cycle, decisive or not, is not braced for trouble. It is simply legible — which, under lights, is the only safety there is.

CASE THIRTY

The Membership Was Always the Electorate: A Gilt Desk and the Wrong Selectorate

A teaching case in political-risk modeling for markets — pricing who actually decides, and holding a view against a loud narrative

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the record's own metadata deflates this call before any critic can: its corpus flag does **not** mark it against-consensus, public Conservative-member polling had shown Truss the stable favourite for weeks before the seal, its published actionability score is a low 44 (the sealed text carried no trading recommendation — the desk decision is this case's construction), and its ~1-in-30 improbability is priced against a diaspora narrative, not the professional prior a rates desk actually holds. The gilt-market crisis, the emergency intervention, and the resignation that followed are **documented aftermath, never foreseen** — the seal named a winner and a mechanism, nothing more. Every ballot figure and market event is real and cited. The **protagonist and his firm are composite characters** — "Oliver Radcliffe" and "Thornbridge Asset Management" are not real; their book and meetings are teaching devices, and no outside-world fact discussed is invented. The case is nonpartisan, nothing in it is investment advice, and it takes **no position** on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

WEDNESDAY, AUGUST 10, 2022 — 07:45, LONDON

Oliver Radcliffe ran UK rates for Thornbridge Asset Management — gilts, sterling rates, the political-risk overlay that both required in a summer when the country was choosing a new Prime Minister. The Conservative leadership contest had reached its final two, and his morning tape ran one story on a loop: **Rishi Sunak led the MPs' ballots.** Every round. The momentum framing wrote itself; the financial press ran it; the diaspora press ran it harder, openly

celebrating the imminent first Indian-origin Prime Minister. Half the political-risk commentary crossing Radcliffe's desk was, in effect, modeling the MPs' ballot as the election.

It wasn't the election. The thing his junior analyst kept saying — with the exasperation of someone repeating the obvious — was that the MPs' rounds were over; the contest had moved to the **roughly 172,000 members of the Conservative Party**, a different body with different politics, whose ballots were already going out, and whose public polling had shown Liz Truss a heavy, stable favourite for weeks. The result was scheduled, dated, certain to arrive: **September 5**.

The item that sharpened the morning was a public, timestamped post of **August 8** — from a source with a sealed and publicly graded record, method disclosed as Vedic astrology — that named the winner and, more usefully, named the *mechanism*: Truss beats Sunak in "a close battle right to the end"; Sunak leads the MPs' ballots, but the membership decides, and the membership — where Truss carried "the establishment and the government machinery" — would hold against a late narrowing.[1] Radcliffe's honest note in the file: *this tells a professionally informed desk nothing the member polls don't.* And then the second, more interesting note: *so why is half my inbound still modeling the MPs' ballot?* The seal's value, if any, wasn't information. It was a dated, falsifiable statement of the **correct selectorate** — a discipline anchor for four weeks of holding an unfashionable base case against the loudest narrative in the market.

THE THREE POSTURES

Package A — Ride the narrative. Let the momentum coverage set the base case; re-price at the September 5 headline like everyone else. Cost: zero now. Exposure: four weeks of political-risk modeling built on the wrong electorate, corrected all at once, late, with the crowd.

Package B — The quiet reweight. Model the actual selectorate: set the desk's base case to the membership outcome (Truss), sized as a directional succession view inside the existing mandate — same-day execution, cost equal to the trade's own spread — and write the one-page selectorate memo that would hold the desk's nerve when the narrative got louder: *who votes, what they poll, what would falsify this* (a genuine member-poll reversal, not another MPs'-ballot story). Nothing in the position went beyond what the seal and the polls jointly supported: *who occupies the office from September 6* — not what its occupant would do there.

Package C — The conviction trade. Size it like information: lever the succession view, take it to clients as a call, cite the full stack including the sealed post. Cost: client-facing reputation staked on a two-horse race the desk's own honest analysis called well-priced — and a source citation no compliance officer would enjoy.

The head of desk's challenge, at the Thursday risk meeting, was the right one: *"Member polling is public. If the membership decides and the polls are stable, the market should already price Truss — where's the trade?"* Radcliffe's answer, half analysis, half confession: *The trade, if there is one, is small and it's against the* narrative*, not the market — positioning drift, client positioning, the commentary flow all lean Sunak-momentum in ways the clean prior doesn't justify. But the bigger P&L isn't this trade. It's the next four weeks of every UK decision this desk makes being modeled on the right electorate. B costs a spread and a memo. The alternative is discovering on September 5 which half of our assumptions were rented from a headline.* What remained for Thursday: the falsifier line, and whether the sealed post belonged in the memo at all.

EXHIBITS

Exhibit 1 — The contest's structure and the tape, early August 2022. Two-round mechanism: MPs' ballots (concluded — Sunak led every round, final 137–113) → postal ballot of ~172,000 party members (in progress, result September 5). Coverage: MPs'-ballot momentum as the live story; diaspora press expectant; member polling showing Truss the stable favourite for weeks.

Exhibit 2 — The sealed input. Public post of August 8, 2022 (on the record): Truss to win "a close battle right to the end"; the mechanism named — Sunak leads MPs, the membership decides, Truss holds against late narrowing. Advisory documentation and grading at jyotishintelligence.com/intel/advisory/IA-UK22-001 — published profile: ~1-in-30 (~4.9 bits, modest against a corpus median near 5.9, and priced against the diaspora narrative rather than the professional prior); SITA 65 at the 52nd percentile, **actionability 44** — no trading recommendation in the sealed text; corpus flag: **not** against-consensus. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The desk's inbound (composite teaching device). Political-risk notes modeling "momentum"; client queries anchored on the MPs' rounds; a book whose UK overlay assumes continuity of the summer's fiscal stance

under either winner — an assumption nobody had priced separately.

Exhibit 4 — What subsequently happened (public record, for the teaching note). September 5: Truss defeats Sunak among members, 81,326–60,399 (57.4–42.6) — "Sunak leads MPs, membership decides" exactly as sealed weeks before either number existed; PM on September 6. **Documented aftermath, never foreseen by the seal:** the September 23 mini-budget of unfunded tax cuts; sterling's record low against the dollar within days; the Bank of England's emergency gilt purchases (September 28) to arrest a disorderly move in long-dated yields; the top-rate reversal (October 3); the Chancellor replaced (October 14); Truss's resignation (October 20) — the shortest premiership in UK history — and Sunak's appointment (October 25).

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Ride the narrative	B · Quiet reweight	C · Conviction trade
What it adds	comfort of the tape	right-selectorate base case, four weeks early	size and noise
What it forecloses	nothing but accuracy	a spread; some meeting friction	client credibility if the close race breaks odd
Cost if Sunak wins	0 (wrong with everyone)	small loss, honest memo	visible loss, sourced embarrassingly
Cost / value if the polls hold	four weeks of mis-modeled UK risk	clean base case into September 6	modest extra P&L, outsized risk
Speed	—	same day	days + client cycle
Authority required	none	desk mandate	CIO + compliance
Reversibility	n/a	full (it's a spread)	reputational drag
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/gilt-desk
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Membership Was Always the Electorate

Rails: the deflations lead — member polling was public and stable, the record's own flag reads *not against-consensus*, the ~4.9 bits are modest and priced against the wrong prior for this reader, and actionability was scored 44 because the desk decision is the case's construction, not the seal's text; the gilt crisis, the Bank's intervention, and the resignation are **aftermath, never foreseen** — any student telling that credits the seal with September 23 onward has broken the case; nonpartisan throughout; nothing herein is investment advice.

LEARNING OBJECTIVES

(1) Selectorate modeling: the first question about any political event is *who actually decides* — and entire narrative ecosystems routinely model the wrong body (MPs for members; polls for electoral colleges; analysts for boards); (2) narrative-versus-prior separation: distinguishing "the market is mispriced" from "the commentary is mispriced" — most political-risk P&L for a desk lives in the second, in the quality of four weeks of downstream decisions, not in the headline trade; (3) the discipline anchor: what a dated, falsifiable statement is worth when its information content is near zero for an informed reader — a written base case with a named falsifier, held against noise (this volume's recurring answer, here in market clothing); (4) honest self-pricing as a source credential: a record that flags its own call as *not* against-consensus, prices it modestly, and scores its own actionability low is exhibiting exactly the properties that make its bigger claims elsewhere readable.

ASSIGNMENT QUESTIONS

1. Write Radcliffe's selectorate memo: base case, sizing, the falsifier line (what member-poll evidence flips the desk), and the citation decision — does the sealed post appear? Defend both the analytic and compliance dimensions of your answer.
2. The head of desk's objection: "if the polls are public and stable, the market prices it — where's the trade?" Answer rigorously: enumerate where narrative-versus-prior gaps show up for a rates desk (commentary flow, client positioning, option skews, the desk's own downstream assumptions), and which of them are monetizable versus merely correctable.
3. Exhibit 3 hides the case's largest exposure: the book assumed fiscal-stance continuity under either winner — an assumption nobody priced separately. The seal said nothing about policy. What *process* (not prophecy) should have surfaced

that assumption in August, and what would examining it have cost?

4. Compare this case with The Exit Polls and The Fifteenth Ballot: three selectorate/consensus failures in three political systems. Extract the common checklist a desk should run on any upcoming political event — five questions, no more.
5. Complete Exhibit 5 and sign it; fill the hindsight row twice — once as of September 6 (the call resolves), once as of October 20 (the aftermath resolves) — and note carefully which entries the seal is entitled to and which it is not.
6. **The machine staff-officer.** An AI political-risk system ingests polls, ballots, and coverage — and, trained on volume, weights the loud MPs'-momentum corpus far above the thin member-poll corpus. Precisely this case's failure, automated. Design the debiasing rule: how should a machine be told to identify the *deciding body* and weight evidence by selectorate relevance rather than by corpus volume — and which check remains human? Write it as a one-paragraph model requirement.

THE BOARD PLAN (75 MIN)

Open on Exhibit 1's mechanism (10): make the class state who votes before any narrative appears — then show them the tape, and watch the momentum story recolonize the room; that recolonization *is* objective 1. The input and its deflations (15): Exhibit 2 read with the record's own modest pricing aloud; run the "what does this add?" question honestly — near zero information, one dated anchor. Q2's gap enumeration (15): where narrative mispricing actually bites a desk. The memo (20): Q1 written, falsifier mandatory, citation debated. Reveal Exhibit 4 (10): the 57.4–42.6 result and the mechanism landing exactly — then the aftermath, read with the rail enforced hard: *none of this was in the seal.* Close (5) with Q4's checklist — the transferable artifact.

WHAT HAPPENED, AND THE GRADE

On September 5, the membership decided, as the mechanism had said it would: Truss 57.4%, Sunak 42.6% — with Sunak having led every MPs' round, 137–113 at the last — and she entered Downing Street on September 6. The record grades the call a **HIT priced modestly and flagged honestly**: ~1-in-30 against the diaspora narrative, *not* marked against-consensus in its own corpus metadata, actionability 44 because the sealed text recommended nothing — the record refusing, in its own file, every inflation a marketer would have reached for. Then came the aftermath the seal never touched and this case never credits it with: the mini-budget, sterling's record low, the Bank of England's emergency gilt operation, the reversals, the resignation — twenty days that repriced UK sovereign risk for a generation of desks. For Radcliffe's composite book, Package B's ledger was as forecast in the memo: a small, clean succession view — and, worth far more, four weeks of UK decisions modeled on the body that was actually voting, entering the most violent quarter in modern gilt-market history with its base assumptions already its own.

EPILOGUE

The library-shelf lesson here is deliberately the quiet one. The seal got the winner; the desk's real prize was a *habit* — who decides, what do they poll, what falsifies — that costs one page per political event and was worth more than any trade the summer offered. Markets meet politics constantly, and the commentary will always model the loudest ballot. The desks that endure are the ones that, every time, ask the junior analyst's exasperated question first: *who is actually voting?*

CASE THIRTY-ONE

The Paralysis Assumption: An Access Plan Waiting for a Window

A teaching case in humanitarian planning — spending scarce staff-hours against an institutional turn the field does not expect

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and this case's underlying call is **ungraded by the record's own policy**: though it clears the record's mechanical bar for a scored hit, the record declines the points, stating plainly that it will not convert a mass-casualty theatre into a scoreboard entry. This case inherits that refusal completely: the human toll of the Gaza war is named and deliberately left unpriced; **no operational effect is claimed for any resolution** — Resolution 2728's demand did not by itself open crossings; and the record's own concessions travel with every page: eventual Council action may have been foreseeable trend rather than foresight, the materialisation was institutional rather than operational, and citing any foresight instrument in such a theatre risks instrumentalising suffering — an objection this case carries rather than deflects. Every date and public event is real and cited. The **protagonist and her cell are composite characters** — "Lena Haugen" and the "Access Planning Cell" are not real; no actual UN body's, consortium's, or donor's deliberations are depicted, and no outside-world fact discussed is invented. The case takes **no position** on the conflict, its parties, or any settlement — and none on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

SUNDAY, NOVEMBER 12, 2023 — 19:00, A COORDINATION HUB IN THE REGION

Lena Haugen led a small access-planning cell inside an INGO consortium's Gaza coordination structure — the function that existed to answer one question whenever the politics moved: *what, exactly, does this sector ask for, of whom, in the hours when asking is possible?* Her resource was brutal in its

scarcity: planning-cell staff-hours, in the middle of the worst weeks any of them had worked. Every hour spent drafting against a hypothetical was an hour not spent on the day's actual convoys, the day's actual denials.

The field's operating assumption that November was as close to consensus as her profession ever came: **Security Council paralysis, through the winter.** The autumn had demonstrated it repeatedly. Serious access planning keyed to a UN institutional turn was, by that assumption, planning for a door that would not open.

The item a colleague had forwarded — with the subject line *"for the scenario shelf, nothing more"* — was a public, timestamped thread of **November 11**, from a source with a sealed and publicly graded record, method disclosed as Vedic astrology. Its fourth statement, verbatim: *"Between January and March 2024, there will likely be a collapse in any remaining system of governance in Palestine. It is during this time that world organizations like the UN will finally muster some courage to attempt playing a role in humanitarian assistance."*[1] The same thread staked a far darker, far longer claim — *"inauspicious changes that will continue for two decades with a lasting impact on the very idea of Palestine as an independent nation state"* — which Haugen read once, filed under a horizon no planner could adjudicate, and set aside.

She held statement four to exactly its own words, because the discipline mattered here more than anywhere: an institutional actor class ("world organizations like the UN"), a humanitarian-assistance activity class, a governance-collapse pairing, a nine-week window. **No Council named. No resolution. No ceasefire demand.** If she costed this scenario, she would be costing that generality and nothing finer. And the costing question was real, because the actions a window assumption unlocked were cheap only *before* any window opened — that was their entire structure.

THE THREE POSTURES

Package A — The paralysis assumption. Spend nothing against the scenario; every staff-hour stays on the day's operations. Cost now: zero, and defensible by the entire field's read. Exposure: if an institutional turn came, the sector would meet its moment of maximum political attention drafting from scratch — asking for generalities in the only week specifics might have been granted.

Package B — Quiet readiness. Three actions, staff-time only, inside existing mandate, startable within days. *One — pre-position the ask list:* the sector's specific, named demands for any institutional turn — crossings, deconfliction

mechanisms, notification systems, fuel — take weeks to agree internally across a consortium and cost nothing but planning-cell time; drafted in winter, they would exist on whatever day the politics moved. *Two — pre-negotiate the language:* operative humanitarian language in any eventual UN product is cheapest to influence before a text exists; briefing the missions and counterparts who would someday shape such language was a standing, within-mandate activity. *Three — re-baseline the counterpart map:* the sealed statement's darker pairing — governance collapse — was, for an access planner, a counterpart-mapping problem: access negotiation presumes an interlocutor, and if governance collapsed the interlocutors would become ad hoc — municipal, clan, commercial, armed. Mapping them was months of work, nearly free to start early, nearly impossible to compress late.

Package C — The pivot. Reorganize the cell around the window: stand down other scenario work, brief leadership that a Q1 turn was coming. Cost: the cell's scarcest resource concentrated on one unverified thread, in a theatre where every misallocated hour had human weight.

The objection came from her own deputy, and it deserved its full force here more than in any case in this book: *"People are dying now. You want winter staff-hours on language for a resolution the entire field says won't come — on the strength of an astrologer's thread? If we're wrong about where to spend these hours, the cost isn't a memo."* Haugen's answer, drafted and redrafted for Tuesday's cell meeting: *Everything in B is work this sector should arguably hold current regardless — an ask list, counterpart maps, relationships with mission drafters. The thread doesn't justify the work; the work justifies itself, and the thread's only role is the date: it is the single input in our stack that prices the paralysis assumption as an assumption. If the field is right, B cost us hours we partly owed anyway. If the field is wrong, the sector meets its one week of leverage with its demands already named.* What she could not draft in advance was the weighting itself — how many hours, from whom, starting when — and that was Tuesday's question.

EXHIBITS

Exhibit 1 — The field's assumption, November 2023. Documented consensus: continued Security Council paralysis through the winter; the autumn's record supporting it. The planning cell's economics: staff-hours as the binding scarce resource, priced in operational time during an active catastrophe.

Exhibit 2 — The sealed input. Public thread of November 11, 2023 (on the record, timestamps platform-fixed): statement four quoted verbatim in the case text — actor class, activity class, governance-collapse pairing, January–

March window; no Council, resolution, or ceasefire demand named. The two-decade statement: open by construction (a ~2043 horizon), displayed as a dated forward call, never banked. Advisory documentation at jyotishintelligence.com/intel/advisory/IA-PS-2023 — **ungraded by the record's own policy**: the institutional call clears the mechanical bar for a hit and the record declines the points, refusing to convert this theatre into a scoreboard entry. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The cell's ledger (composite teaching device). Ask list: fragments, unagreed across the consortium. Counterpart map: current for pre-war structures only. Mission relationships: strong with two delegations, thin elsewhere. Available hours: the honest number — small, and every one contested.

Exhibit 4 — What subsequently happened (public record, for the teaching note). March 25, 2024 — 135 days after the seal, in the final week of the stated window: the Security Council passes **Resolution 2728**, its first demand for an immediate ceasefire in Gaza, the United States abstaining. The demand did not by itself open crossings, and no operational effect is claimed. The two-decade statement remains open; the record's timeline logs one dated development — the February 2025 displacement-premised "Gaza Riviera" proposal — as directionally consistent, and nothing more.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Paralysis assumption	B · Quiet readiness	C · The pivot
What it adds	every hour to today's operations	asks named, language seeded, counterparts mapped — before any window	total focus on one scenario
What it forecloses	the cheap version of readiness, which exists only before the turn	some winter hours, honestly contested	the rest of the scenario shelf
Cost if paralysis holds	0	hours partly owed anyway	hours taken from live needs, on one thread
Cost / value if a window opens	drafting during the only week of leverage	demands ready in the week attention peaks	ready — at a cost no one can defend
Speed	—	days to start, weeks to mature	immediate, wrenching

Authority required	none	cell's own mandate	consortium leadership
Reversibility	n/a	full	organizationally painful
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading — and this record refused to grade)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/access-plan
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Paralysis Assumption

Rails — read aloud, and enforce without exception: the human catastrophe is named and never priced — the record refuses, as policy, to convert this theatre's toll into any metric, and every classroom telling inherits that refusal; **no operational effect is claimed for Resolution 2728** — the materialisation was institutional, and crossings are a different fact; the foreseeable-trend objection is conceded (months into any war, eventual Council action has a rising base rate); the instrumentalisation objection is carried, not deflected — discussing foresight instruments in a mass-casualty theatre is itself ethically loaded, and the record's refusal to grade is the honest posture the class should interrogate; the case takes no position on the conflict or its parties; the two-decade statement belongs to a scenario shelf, never a forecast column.

LEARNING OBJECTIVES

(1) Assumptions are budget lines: "the Council is paralyzed" was not a belief but an allocation — zero hours to window-readiness — and the case teaches pricing standing assumptions as the spending decisions they are; (2) the before/after cost structure of institutional turns: when politics move, instruments are drafted in days — so the sector's asks, language, and relationships are cheap only in advance, and readiness has an expiry structure the paralysis assumption hides; (3) work that justifies itself: the deputy's objection is met not by defending the thread but by locating the subset of actions owed under any assumption — and being honest about which hours B adds beyond that subset; (4) the ethics of refusal: what it means that a record *declined* a scoreable hit here — students should debate whether the refusal is integrity, marketing, or both, with the record's own concession (an ungraded "hit" is the record's characterisation, never externally adjudicated) in full view.

ASSIGNMENT QUESTIONS

1. Write Haugen's Tuesday allocation memo: which hours, whose, starting when — with the deputy's objection quoted in full at the top and answered underneath, honestly.
2. The foreseeable-trend objection: construct the base-rate case that Council action by spring was likely regardless. What, precisely, did the sealed window add over that base rate for *this* desk — and is "a date that prices the assumption" a real answer or a rhetorical one? Argue both sides.
3. Map the counterpart problem: if governance collapses, list the interlocutor classes an access negotiator inherits, what mapping each requires, and why the exercise compresses so badly under time pressure. (This question is answerable entirely

without the forecast — which is objective 3 made concrete.)

4. The record refused to grade this call. Debate the refusal: what does a scoreboard-decline protect, what does it still quietly claim, and what adjudication standard *would* satisfy a skeptic in theatres like this one? Connect to the City Hall Window's declined hit — same discipline, opposite valence.
5. Complete Exhibit 5 and sign it; fill the hindsight row as of March 26, 2024, writing within the rails: the institutional turn came in-window; the crossings are a different fact; the human toll is not yours to price.
6. **The machine staff-officer.** An AI scenario system can maintain the ask list, the language bank, and the counterpart map continuously, for every theatre at once — making Package B's marginal cost near zero and the paralysis assumption's hidden budget line visible everywhere. What does that change about this case's dilemma, what does it not change (consortium agreement is human-speed; mission relationships are human-held), and what special governance does AI-assisted planning owe a theatre where the record itself refuses metrics? Write the cell's one-paragraph policy.

THE BOARD PLAN (90 MIN — THIS CASE AND THE WARNING DELIVERED TWICE CARRY THE VOLUME'S HEAVIEST SETTINGS; SCHEDULE ACCORDINGLY)

Open on the cell's economics (15): staff-hours as the scarce resource, priced in operational time — the deputy's objection should be voiced early and taken seriously; this case fails if the room treats it as an obstacle rather than the strongest position on the board. The assumption-as-budget-line insight (15): objective 1, built from Exhibit 1. The input, bounded hard (15): Exhibit 2 at its own generality — run the sealed-text discipline visibly; anything Council-specific is outcome, not input. Package B and the owed-anyway analysis (20): Q3 worked without the forecast, then the honest delta named. Commit (10): memo + Exhibit 5. Reveal Exhibit 4 (10): Resolution 2728, 135 days, final week of the window — and the rail, restated: institutional, not operational. Close (5) with Q4's refusal debate, left deliberately unresolved.

WHAT HAPPENED, AND THE RECORD'S REFUSAL

On March 25, 2024 — 135 days after the seal, in the last week of the stated January–March window — the Security Council passed Resolution 2728, its first demand for an immediate ceasefire in Gaza, with the United States abstaining. The consensus paralysis assumption broke on a date inside the sealed window. The demand did not by itself open crossings, and this case claims no operational effect for it — the materialisation was institutional, exactly as the record characterises it. And then the record did the thing this case exists to examine: **it declined to grade its own call.** By its published mechanics the statement clears the bar for a scored hit; the record refuses the points, stating that it will not convert a mass-casualty theatre into a scoreboard entry — so no bits are banked, no Brier term contributed, and the "hit"

status remains the record's own characterisation, never externally adjudicated, a concession the file prints about itself. For the composite cell, the postures priced as their structure dictated: a sector that had named its asks in winter met the week of 2728 with demands on paper; a sector that had assumed paralysis met it drafting — and either way, the crossings the next morning answered to facts no resolution, forecast, or planning cell controlled.

EPILOGUE

This case sits beside The Warning Delivered Twice at the volume's gravest edge, and together they mark the boundary the whole book must respect: decision discipline is teachable everywhere, but not everything that resolves is a result, and not every result belongs on a scoreboard. The residue for a planner is the case's quietest line — *an assumption is a budget; price it like one* — and the residue for the profession that keeps records is the refusal itself: kept misses and declined credit are the same discipline, facing opposite directions, and a field that learns to demand both from its forecasters will deserve the instruments it gets.

CASE THIRTY-TWO

Finishing the File Before the Seam: The Succession Contingency

A teaching case in policy planning — converting a standing worry into dated work, on the file no desk ever finishes

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and this case is bound by the record's own stated constraints, which it holds absolutely: **it claims nothing about His Holiness the 14th Dalai Lama — not his health, not his longevity, not that any named date means anything.** His Holiness's own stated expectation — that he may live to 130, restated at his 91st birthday in Leh — is carried at full weight; the record **excludes this advisory from its graded ledger on ethics** and states, on this file uniquely, that it would rather be wrong. The dated windows function in this case as exactly one thing: a completion deadline attached to contingency work the desk already owed on structural grounds that predate and do not depend on any forecast. The **protagonist and his ministry are composite characters** — "Thomas Beck" and his policy-planning staff are not real; no actual government's deliberations are depicted, and no outside-world fact discussed is invented. The case takes **no side** on the contested question of recognition, and **no position** on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

MONDAY, SEPTEMBER 22, 2025 — 11:00, AN ALLIED FOREIGN MINISTRY

Thomas Beck owned, among a dozen contingency files on a policy-planning staff, the one every serious Indo-Pacific desk carries and almost none has finished: *what do we do when the succession of His Holiness the 14th Dalai Lama opens?* The file's structural facts were old, public, and independent of anyone's forecast. Beijing asserts the authority to approve the reincarnation — State Religious Affairs Bureau Order No. 5, in force since 2007 — and has run the play before, taking the six-year-old 11th Panchen Lama in 1995 and installing its own candidate. A contested succession was therefore close to a structural

certainty rather than a tail risk; the danger was the **interregnum** — the gap between a transition, whenever it came, and a recognized successor, in which a state actor could move at the speed of a prepared bureaucracy while democracies moved at the speed of an interagency process.

And yet the file sat, as such files do, perpetually unfinished — because it had **no deadline**. Work without a date loses every quarter's triage to work with one. His Holiness was ninety, in apparent good health, and had told the world at his birthday celebrations he might live to 130 — a statement Beck's file carried, as it should, at face value and full weight.

The item that forced the question was a public advisory of **September 18** — a sealed video and thread from a source with a publicly graded record, method disclosed as Vedic astrology — that named four dated windows between February 2, 2026 and February 5, 2027, and counselled that succession-contingency arrangements be **completed before the first of them**. Beck read the record's own framing twice, because it was unlike anything else in his intake: the publisher had *excluded the advisory from its own scoring ledger on ethics* — tracked, never graded — declared it was not a statement about anyone's health, and written that on this file it would rather be wrong.

Beck's reaction split cleanly in two, and he wrote both halves down. **As forecast, he refused it entirely** — his government did not, would not, and should not date the life of a revered living person, and the advisory's four-windows-plus-envelope structure was, as his own China analyst put it, close to unfalsifiable anyway. **As a deadline, it solved his oldest problem.** Every element of the file's owed work — recognition language, allied coordination, sanctions folders under statutes that had existed since December 2020, an interregnum communications plan, an engagement protocol with the Gaden Phodrang — was justified by Order No. 5 and the Panchen precedent alone, datelessly, and had been for years. What the work had never had was a *completion date* that could win a triage meeting. The question he took to Thursday's planning review was narrow and strange: **may a planning staff take a deadline from an input it refuses as a forecast — and if not from this, then from what?**

THE THREE POSTURES

Package A — The standing file. Keep the work undated and owed-someday, as every peer desk did. Cost: zero, and the company of the entire field. Exposure: the interregnum arriving — whenever it arrived, this year or in decades —

with the language undrafted, the allies unaligned, and every first-week choice improvised against a prepared counterparty.

Package B — Take the deadline, refuse the forecast. Convert the file into a dated completion program — target: complete before February 2, 2026, 133 days out — while stating in the program memo itself that the date was a planning device, not a prediction, and that the work's justification was structural and forecast-independent. The elements: *pre-agreed recognition language* (what the government will and will not say about a Dharamsala-lineage successor and about any state-installed candidate); *allied coordination*, India above all — the venue of any lineage recognition, the partner whose equities (the border, the exile community, Beijing's pre-declared position) made its first-week posture the least improvisable and the slowest to align, and therefore the single strongest argument for starting against a clock; *sanctions folders* under existing statutory authority — target identification, no new law needed; *the interregnum communications plan*, built for an information environment already contested; *the Gaden Phodrang engagement protocol*. Staff work throughout — weeks, memoranda, consultations; no cabinet decision required.

Package C — The visible posture. Brief upward that a transition window had been forecast; elevate the file politically; open allied conversations framed by the dates. Cost: everything the reverence constraint existed to prevent — a government trafficking in dated speculation about a living spiritual leader, the offense to the community concerned, and the diplomatic gift of the leak.

The deputy political director's objection at Thursday's review was the one Beck respected most: *"If the work is justified without the dates — and it is — then use the real justification and set our own deadline. Taking a completion date from an astrologer's window on a revered man's life is ethically fraught even in a classified annex, and if it ever surfaced, indefensible."* Beck's draft response, which he was still testing when the review convened: *You're right about the justification, and the memo will say so — Order No. 5 and 1995 are the reasons; no forecast appears in the rationale. But notice what our own history shows: we've had the real justification for eighteen years and never finished the file, because self-imposed deadlines on someday-files lose every triage fight to dated work. The window's only function is that it is a date we did not choose, attached to a completion target we should have had anyway. If we can substitute an internal date with the same binding force — genuinely binding, surviving the next three triage cycles — I'll take*

ours over theirs without a second thought. Can we? That question — whether the ministry could bind itself without borrowing a clock — was the one the review had to answer.

EXHIBITS

Exhibit 1 — The structural file (forecast-independent, decades old). Order No. 5 (2007): state approval asserted over reincarnation. The 1995 precedent: the 11th Panchen Lama, taken at six, replaced by a state candidate. US statutory authority in force since December 2020 (P.L. 116-260 Div. FF): succession decisions belong to Tibetan Buddhists; sanctions available for interference. The interregnum as the decisive period; first-week statements as the contest's terrain. His Holiness's stated expectation: he may live to 130 — carried at full weight.

Exhibit 2 — The sealed input, and the record's refusal. Advisory of September 18, 2025 (sealed video + thread, on the record): four windows, February 2, 2026 – February 5, 2027; counsel to complete arrangements before the first. The record's own terms: **excluded from graded scoring on ethics**; no statement on health; explicit that on this file it would rather be wrong. Documentation at jyotishintelligence.com/intel/advisory/IA-DL-2025. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The desk's history with the file (composite teaching device). Opened: years ago. Elements drafted: fragments. Triage record: deferred in each of the last six quarterly cycles, each time to dated work. Peer desks: same posture, per liaison.

Exhibit 4 — What the following year held (public record, for the teaching note; context, never confirmation). Early February 2026: His Holiness the target of a distorted files-based controversy — firmly denied on official letterhead, independent reviews finding no evidence of wrongdoing; contemporaneous analysis framing it as a live information war against the institution. July 9, 2026: TIME reports Beijing preparing to assert that the next reincarnation "must be approved by the central government"; mainstream analysis now anticipates two rival Dalai Lamas; the July 27 cover: "HIS LAST STAND." June 2026: a congressional amendment moves to place succession beyond CCP jurisdiction. July 1, 2026: new religious-affairs assimilation measures take effect. The question the advisory told planners to finish

preparing for became, within ten months, the front page — while the file, on most desks, stayed open and undated. **Two sealed windows remain ahead as of this writing; nothing is claimed about them.**

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Standing file	B · Deadline taken, forecast refused	C · Visible posture
What it adds	nothing	a completion date that survives triage	political attention
What it forecloses	nothing — perpetually	133 days of staff time, quietly	the reverence constraint, publicly
Cost if the windows mean nothing	0	a finished file the desk owed anyway	offense + leak exposure, for nothing
Cost / value whenever the interregnum comes	improvisation against a prepared actor	language ready, allies aligned, folders built	prepared — at an indefensible provenance
Speed	—	weeks, inside existing authority	political calendar
Authority required	none	planning staff's own	ministers
Reversibility	n/a	full (the work keeps)	poor
Recommendation & rationale		*(yours)*	
In hindsight	*(this record refuses to grade this file — write yours anyway)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/succession-file
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — Finishing the File Before the Seam

Rails — absolute, and read aloud: the case, like the record, **claims nothing about His Holiness** — health, longevity, or the meaning of any date; his stated expectation of living to 130 stands at full weight in every telling; the record's ethics exclusion is honored — no student telling scores the windows, speculates about them, or treats Exhibit 4 as confirmation of anything; the deputy's objection is the case's strongest voice and may win; the reverence owed a living spiritual leader and the community around him governs the room's tone throughout; no side is taken on recognition.

LEARNING OBJECTIVES

(1) The deadline posture at its purest: someday-files lose every triage cycle to dated work — so the scarcest planning commodity is not analysis but a *completion date with binding force*, and the case asks where such dates may legitimately come from; (2) justification hygiene: separating a program's *rationale* (structural, forecast-independent, publishable) from its *scheduling trigger* (possibly neither) — and writing the memo so the two never contaminate; (3) the interregnum as a general class: successions, transitions, and gaps between authority are decided by pre-agreed language and pre-aligned allies, at press-office speed, against prepared counterparties — the file's value concentrates entirely before the seam; (4) ethical intake boundaries: some inputs carry constraints heavier than accuracy — and a record that *refuses to score itself* on such a file models the boundary the profession should demand.

ASSIGNMENT QUESTIONS

1. Write Beck's program memo under Package B: the structural rationale (no forecast appears), the completion date, the workstreams and owners — then write the classified annex's one honest paragraph about the scheduling trigger, and defend both documents surviving a leak.
2. The deputy's challenge is the case's hinge: *can* a ministry bind itself to a self-imposed deadline with real triage force? Propose the mechanism (a minister's signature? a standing review? external anchoring?) — and test it against Exhibit 3's six deferrals.
3. The India paragraph: why is the slowest, hardest coordination the strongest argument for starting earliest? Generalize — in any contingency file, identify the element whose lead time alone justifies the whole program's clock.

4. Compare this case with The Founder Window: both take deadlines from windows on a person. What makes one a calendar entry and the other an ethical problem? Write the boundary rule a planning staff should adopt for person-linked windows — and apply it to both cases honestly.
5. Complete Exhibit 5 and sign it. The hindsight row cannot be graded — the record refuses, and two windows remain open. Write instead what *completion* would look like, dated, and what the desk owes the file regardless of every window passing empty.
6. **The machine staff-officer.** An AI planning system can draft recognition language, maintain sanctions folders, and track every allied position — making Package B's marginal cost near zero and removing the triage excuse entirely. It can also ingest person-linked windows without any ethical hesitation at all. Which half of this file may the machine hold, which half must it never touch (the scheduling trigger? the Gaden Phodrang protocol?), and who writes the boundary? Draft the planning staff's rule for AI on reverence-constrained files.

THE BOARD PLAN (90 MIN — TONE MATTERS HERE AS MUCH AS CONTENT)

Open on Exhibit 1 only (20): build the structural case — Order No. 5, 1995, the interregnum, the statute — until the room agrees the work is owed with no forecast anywhere in sight; then confront Exhibit 3: *why is it unfinished everywhere?* The someday-file problem is objective 1, and it must be felt before the input appears. The input, with its refusals (15): Exhibit 2 read together with the record's ethics exclusion and its would-rather-be-wrong line; let the room sit with how unusual that posture is. The deputy's objection (20): Q2 argued fully — the self-binding question is the intellectual center, and "we cannot, actually" is an admissible finding. Commit (15): Q1's two documents + Exhibit 5. Exhibit 4, read as context only (10): the information war, the front pages, the amendment — with the rail restated: nothing here confirms anything. Close (10) with Q4's boundary rule, drafted by the room.

WHAT HAPPENED — AND WHAT THIS CASE REFUSES TO SAY

Within ten months of the seal, the succession question moved from standing file to front page: Beijing's preparation to assert approval authority reported by TIME, rival-Dalai-Lamas analysis mainstream, a congressional amendment moving, assimilation measures in force — the world's press, legislatures, and the counterparty's own regulatory apparatus all now treating as imminent-horizon work what most contingency files still hold undated. His Holiness, at his 91st birthday in Leh, restated that he may live to 130 — and that statement stands, in this case as in the record, at full weight. **Two sealed windows remain ahead; nothing is claimed about them; the record tracks and will not score; and this note grades only the planning question:** a desk that took the deadline finished, in 133 days of ordinary staff work, a file whose structural justification had been sufficient for eighteen years — and a desk that

refused the provenance owed itself, by the deputy's own argument, a binding date from somewhere. The case's finding is that "somewhere" is the hard part, and always was.

EPILOGUE

Every planning staff carries a drawer of someday-files — successions, transitions, ruptures — each justified beyond argument and none funded past the next triage. This case, the most constrained in the volume, exists for that drawer: *the rationale you publish must stand without the clock; the clock may come from anywhere you can defend — including from yourself, if you can make yourself believe it.* And where a file touches a living person owed reverence, the record that sealed the input has already drawn the line the profession should keep: track the work, refuse the score, and prefer, openly, to be wrong.

CASE THIRTY-THREE

Fifty-Three Days Before Norris: The Basin-Access Decision

A teaching case in natural-hazard management — cheap reversible posture against a window whose headline you should refuse

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and this case's sealed headline **missed outright**: no eruption occurred, the U.S. Geological Survey held the Volcano Alert Level at NORMAL throughout the window, and **exact-time eruption forecasting has no scientific basis** — a fact this case states in its own voice, not as a concession extracted from it. The in-window event that did occur (the April 15 Norris hydrothermal explosion) was small and produced no reported injury, so the counterfactual value of any earlier posture change **cannot be shown to have prevented anything**. The advisory is excluded from the record's graded ledger as prevention-class. Every geological and monitoring fact is real and cited; USGS's NORMAL posture was, on its own instrumented terms, correct throughout. The **protagonist is a composite character** — "Dr. Erin Kovac" is not a real person; no actual National Park Service or USGS staff deliberations are depicted, and no outside-world fact discussed is invented. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology**. The case treats the record as a dated document and leaves the epistemic question to the classroom.

WEDNESDAY, MARCH 27, 2024 — 08:30, A HAZARDS-PROGRAM OFFICE IN THE NORTHERN ROCKIES

Dr. Erin Kovac managed hydrothermal-hazard posture for the thermal basins of a caldera park — the boardwalk closures, the signage tiers, the ranger briefings that stood between four million annual visitors and ground that could, without meaningful warning, turn explosive. Her program ran on two settled professional facts. **First:** the standing hazard was real and permanent — hydrothermal explosions are Yellowstone-class systems' signature surprise, and the USGS's own specialists put it plainly: a catastrophic one "could very well happen today." **Second:** the hazard was *undatable* — the scientific

consensus, backed by the instruments her partners at the volcano observatory ran, was that no method existed to forecast an eruption's date, and the Volcano Alert Level stood at NORMAL.

The item a colleague had forwarded overnight violated the second fact as hard as an input could. A public, timestamped post of **March 26** — from a source with a sealed and publicly graded record, method disclosed as Vedic astrology — named a dated hazard window for the caldera: rising seismicity from April 1; a headline eruption on **May 18, around 2 p.m., ±3 days**; and, separately, in the post's own words, *"at a bare minimum, the possibility for significant hydrothermal explosions in the caldera is extremely high during April and May."*[1]

Kovac's scientific judgment on the headline took one sentence, and she wrote it in the file without hedging: *a dated eruption call has no basis in any science my program recognizes, and I refuse it as such.* What kept the item out of the bin was the professional discipline this book has been teaching from its first case: the input was not one claim but three — a mechanism narrative (seismicity cadence), a headline she refused (the dated eruption), and a **floor clause** (hydrothermal explosions, elevated, April–May) that named exactly the hazard her program owned, in a window, at a class level. Her question was never whether to believe any of it. It was the one that had run through twenty-four prior cases and belonged, if anywhere, here: *what does the floor clause make cheap — and what would it cost me to act as if it might matter?*

THE THREE POSTURES

Package A — Monitor. File it; run the standing posture; trust NORMAL, which her own partners' instruments justified. Cost: zero. Exposure: honestly small — and that fact belongs in the analysis rather than around it.

Package B — Quiet readiness. Three actions, hours of staff work, inside her program's existing authority, all reversible on a stand-down date she would set herself (May 31). *One:* **sharpened posture at the highest-energy features** — supplemental signage, a tightened patrol interval, and a pre-drafted boardwalk-segment closure order held ready. Paperwork and staffing, not construction. *Two:* **a ranger briefing keyed to the window** — front-line staff at the highest-risk basins briefed on an elevated-hazard posture for a defined nine-week span, one line added to briefings they already received. *Three:* **a monitoring-continuity check** — confirming, ahead of the window, that the seismic and deformation network her program depended on had no known

gaps in power or data continuity for the nine weeks flagged; a systems check justified independently of anyone's forecast. None of the three required the eruption date to be credible; each was priced against the floor clause alone.

Package C — Elevated posture. Basin closures, public communication of a hazard window, visitor-management changes. Cost: real operational and reputational money spent against a single outside post, in direct public contradiction of the observatory's NORMAL — a posture that would damage the program's credibility with its scientific partners whether or not anything happened.

Her deputy's objection carried the observatory's full weight: *"USGS has instruments in the ground and holds NORMAL. The eruption clause is pseudo-science on its face — you've said so yourself. If we sharpen posture on an astrologer's window and it leaks, we've told the public the instruments don't govern our decisions."* Kovac's draft reply, for Friday's program meeting: *The eruption clause is refused — it appears nowhere in anything I sign. What I'm pricing is a class-level floor clause about a hazard we describe, in our own literature, as 'could happen today, undatable.' For a hazard like that, a nine-week sharpened posture with a stand-down date is almost the definition of a cheap hedge — and everything in B is defensible as routine vigilance under our own doctrine, no window required. The honest question isn't the source. It's whether B's actions are worth doing on our own logic — and if they are, whether it should bother me that a strange input was what put them on this week's agenda instead of someday's.* The meeting would decide — and the leak question deserved a real answer before it.

EXHIBITS

Exhibit 1 — The professional ground truth. The standing hazard: hydrothermal explosions as the system's signature risk — "could very well happen today" (USGS specialist assessment); undatable by scientific consensus; Volcano Alert Level NORMAL throughout the period; the observatory's instrumented monitoring as the program's governing input.

Exhibit 2 — The sealed input. Public post of March 26, 2024 (on the record): the seismicity narrative; the dated eruption headline (May 18 ~2 p.m. ±3 days) — **which missed outright**; the floor clause on hydrothermal explosions, "extremely high during April and May." Advisory documentation at jyotishintelligence.com/intel/advisory/IA-YS-2024 — headline recorded as MISS; the asset **excluded from the graded ledger as prevention-class** (a risk-watch that does not fire is a counterfactual; the file notes that grading it as a miss would *lower* the ledger's Brier average, so exclusion is not score-

management); the floor clause priced at ~1-in-150 (~7.2 bits) with **no credit taken for the failed eruption date**. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The program's authority map (composite teaching device). Inside her signature: signage tiers, patrol intervals, ranger briefings, pre-drafted closure orders, monitoring-continuity checks. Requiring superintendent or partner concurrence: basin closures, public hazard communications, anything contradicting the observatory's posture.

Exhibit 4 — What the window held (public record, for the teaching note). No eruption; NORMAL throughout; May seismicity at background (62 located events, max M1.8). April 15, 2024: a hydrothermal explosion at Norris Geyser Basin — **the first instrumentally detected hydrothermal explosion in Yellowstone's recorded history** (USGS), leaving a 1–2 m crater on Porcelain Terrace; small; no reported injuries; inside the named window and matching the floor clause. July 23, 2024 — *outside* the window, with no precursors: a larger hydrothermal explosion at Biscuit Basin. Context in the same spring: the strongest space weather in ~21 years (a G5 storm May 10–11; an X8.7 flare May 14), independently forecast by NOAA days ahead — context only, claimed for nothing.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · Quiet readiness	C · Elevated posture
What it adds	nothing	sharpened signage + briefed rangers + verified monitoring, all with a stand-down date	visible protection
What it forecloses	nothing measurable	hours of staff time	credibility with the observatory and the public
Cost if the window is quiet	0	a stand-down memo on May 31	a leaked absurdity, instruments overruled
Cost / value if a thermal event lands in-window	posture as on any other day — likely adequate	posture already sharpened where it happened to matter	protection nobody could justify ex ante
Speed	—	days	weeks + concurrences

Authority required	none	her own	superintendent + partners
Reversibility	n/a	full, on her own date	slow, public
Recommendation & rationale		*(yours)*	
In hindsight	*(this asset is excluded from grading — write yours anyway)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/thermal-basin
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — Fifty-Three Days Before Norris

Rails: the headline **missed, outright, and is said first** — no eruption, NORMAL throughout, and dated eruption forecasting has no scientific basis (the case's own voice, not a concession); Norris was **small, injury-free**, and would very likely have been managed identically under any posture — "might have" is the ceiling and mostly isn't even claimed; USGS's instrumented judgment was correct on its own terms and is treated with full respect; the Biscuit Basin explosion — larger, outside the window, precursor-free — is the standing-hazard reality check that cuts against *everyone's* windows; the prevention-class exclusion is examined, not celebrated.

LEARNING OBJECTIVES

(1) Clause surgery under scientific refusal: how to refuse a headline on disciplinary grounds — explicitly, in writing — while still pricing a class-level floor clause the desk's own doctrine already endorses; the sharpest version of this volume's decomposition discipline; (2) the undatable-hazard paradox: when a hazard is officially "could happen today," every dated window is scientifically empty *and* every cheap sharpening is doctrinally justified — so the real function of any window, credible or not, is agenda-setting for vigilance that never wins the someday-queue (the Festival Window's cousin, in geology); (3) authority-map discipline: Package B lives entirely inside one signature, and Package C's real cost is institutional trust with a scientific partner — students learn to draw the map before pricing the postures; (4) grading edge-cases: an asset whose headline missed, whose floor clause landed on a first-in-recorded-history event, and which the record *excluded* from scoring — the class adjudicates whether the exclusion logic (grading it would have *helped* the average) holds up.

ASSIGNMENT QUESTIONS

1. Write Kovac's Friday program memo: the refusal sentence (verbatim, signed), the three actions justified on the program's own doctrine, the stand-down date, and the leak answer — what does she say when a reporter asks whether an astrologer set the park's posture?
2. The deputy's instruments-versus-inputs argument deserves rigor: what is the proper relationship between an observatory's alert level (instrumented, mechanism-based) and an access-manager's posture (risk-tolerance-based)? Show

they answer different questions — and where NORMAL does and does not bind Kovac.

3. Norris: first instrumentally detected hydrothermal explosion in the system's recorded history, inside the window, matching the floor clause — small, harmless. Biscuit Basin: larger, outside every window, no precursors. Write the two-paragraph honest analysis of what this pair does and does not license about the input, using the record's own ~1-in-150 floor-clause pricing and its refusal of the eruption date.
4. The prevention-class exclusion: the record argues a risk-watch that doesn't fire can't be scored, and notes that scoring this one would have *improved* its average. Adjudicate — is that reasoning sound, self-serving, both? Compare with case 15's prevention-class refusal and case 18's declined hit; draft the general exclusion rule you would hold this record to.
5. Complete Exhibit 5 and sign it; fill the hindsight row as of June 1 (stand-down day) and note honestly how little the world let you verify.
6. **The machine staff-officer.** An AI hazards assistant ingests every input class — instruments, base rates, and sealed windows alike — and proposes posture sharpenings with stand-down dates automatically, at near-zero marginal cost. The someday-queue problem vanishes; so does the human act of *refusing a clause in writing*. What is lost if the refusal sentence disappears from the file — and write the rule that keeps disciplinary refusals explicit even when a machine prices every hedge as cheap.

THE BOARD PLAN (75 MIN)

Open on Exhibit 1's paradox (15): "could happen today" + "undatable" — have the class feel how this pair makes vigilance both permanently justified and permanently deferrable; objective 2's distinction (alert level vs. access posture) should emerge here. The input and the surgery (15): Exhibit 2; the refusal sentence written on the board *first*, then the floor clause priced separately — the order is the discipline. The authority map and packages (15): Q2's boundary, Package C's real cost named as partner trust. Commit (15): memo + Exhibit 5, leak answer mandatory. Reveal Exhibit 4 (10): NORMAL vindicated, the headline dead, Norris first-in-history and harmless, Biscuit Basin outside everything — the full untidy truth, resisting every neat narrative. Close (5) with Q4's exclusion debate, and the volume's running theme stated once more: the record's treatment of its own failures is the only reason any of its windows deserved an hour.

WHAT HAPPENED, AND THE NON-GRADE

No eruption occurred. The Volcano Alert Level stayed NORMAL from first day to last; May's seismicity was background; the dated headline — the clause with the hour attached — missed as completely as a claim can miss, exactly as the scientific consensus on eruption forecasting predicts such claims must. Inside the same

window, on April 15, Norris Geyser Basin produced the first instrumentally detected hydrothermal explosion in the system's recorded history — a 1–2 meter crater, no injuries, squarely in the floor clause's class — and on July 23, outside every window and without precursors, Biscuit Basin produced a larger one, as if to sign the standing-hazard doctrine personally. The record holds the asset **outside its graded ledger** as prevention-class, prices only the floor clause (~7.2 bits), takes no credit for the failed date, and notes that grading the headline would have flattered its average — an exclusion this note asks students to adjudicate rather than accept. For the composite program, Package B's ledger was a few staff-hours, a briefing line, a verified network, and a stand-down memo dated May 31 — a hedge the world never tested, against a hazard that remains, today as always, undatable and real.

EPILOGUE

This is the volume's only natural-hazard case, and it earns its place by inverting the book's usual arithmetic: here the *headline* deserved refusal on first principles, the *hedge* was justified without any forecast at all, and the window's whole contribution was moving never-scheduled vigilance onto one week's agenda. The residue for hazard professionals everywhere — geological, cyber, financial: *for hazards your own doctrine calls permanent and undatable, the cheap sharpening is always defensible on your own logic; write the refusal, take the agenda-setting, set the stand-down date, and let the instruments keep governing.* And the residue for the record-keepers: a window that cannot fail has no business being scored — and saying so, out loud, is the price of being read at all.

CASE THIRTY-FOUR

Planning on Continuity: The Grain Ledger

A teaching case in food-security analysis — choosing the scenario the world balance sheet is built on

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the two strongest objections travel at full strength from the first page: Russia entered the war as the **world's number-one wheat exporter with the record 2022 harvest largely in the bins by the seal date** (a hostile reader calls "remain strong" a momentum call with a hedged extension), and **food and fertilizer were deliberately carved out of the Western sanctions architecture** — so continuity may measure sanction *design* rather than any divergent read. A counter-receipt from the same sealed sitting is printed as the record prints it: the resources clause met a mixed outcome (2023 oil-and-gas budget revenues fell 23.9% from the 2022 peak). Every harvest, export, and delivery figure is real, cited, and kept in its exact units. The **protagonist and her unit are composite characters** — "Dr. Amina Touré" and the "Global Supply Outlook Unit" are not real; no actual agency's deliberations are depicted, and no outside-world fact discussed is invented. The case claims **no downstream welfare outcome** beyond documented flows, and takes **no position** on sanctions policy or on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

MONDAY, NOVEMBER 21, 2022 — 09:30, A FOOD-SECURITY AGENCY

Dr. Amina Touré ran the Global Supply Outlook Unit — the analytic cell whose product, a world grain balance sheet, anchored decisions far heavier than itself: donor-ministry appeal sizing, import-dependent countries' precautionary reserves, the diplomatic capital spent courting alternative suppliers. Her unit's product for 2023–24 was due for its scenario lock, and the scenario question of the season had one name: **Russia**.

The working fear across her field was specific and operational, not academic. Payment friction, shipping and insurance barriers, and fertilizer-trade disruption were expected to reach the Russian food economy and, through it,

the world's — that fear was why the Black Sea grain diplomacy existed, why the fertilizer carve-outs had been negotiated at all. Her unit's draft balance sheet, like her peers', carried a **disruption-weighted** Russian supply case, with analytic effort flowing into stress tests of the Russian production base.

The input her deputy flagged — *"against the entire room, for what it's worth"* — was tweet 4 of a 23-part sealed thread posted **November 19**, from a source with a public, graded record, method disclosed as Vedic astrology: *"Agriculture and food supply will remain strong in Russia. They might even export these to other countries for profit."*[1] Two sentences. Continuity, and continuity weaponizable.

Touré's first honesty was the deflationary one, and she wrote it into the file herself: Russia had entered the war as the world's top wheat exporter; the record 2022 harvest was substantially gathered by the seal date; and food was deliberately outside the sanctions net. On those facts, "remain strong" read like momentum, and a hostile reviewer would say so. Her second honesty cut the other way: *if* those facts made continuity the better base case, then her own unit's disruption weighting — and her whole field's — was mispriced against information everyone already had, and the strange post was merely the only dated artifact in her intake willing to say it plainly. The scenario lock was Thursday. The question was not whether to believe an astrologer. It was which case anchored the balance sheet that everything downstream would inherit.

THE THREE POSTURES

Package A — Hold the disruption weighting. Keep the consensus case; let outside confirmation force any rewrite. Cost: zero, and the field's company. Exposure: a supply-side base case wrong in the *expensive* direction — reserves and appeals sized to a strangulation that never came, analysis spent on a production collapse instead of the constraints that actually bound, and the analytic predicate for Russian grain-as-statecraft unthinkable inside a strangled-sector frame.

Package B — Quiet reweighting. A memo, executable this week, inside her existing analytic authority: anchor the 2023–24 balance sheet on **Russian continuity at or above trend**; redirect the freed stress-testing effort to the genuinely binding constraints (routing, insurance, buyer-side payment channels); flag to reserve- and appeal-sizing clients that the disruption premium in their tails deserved review; and open one scenario page her

field's frame had made unthinkable — *a strong, sanctions-immune food sector used as a foreign-policy instrument* — at class level, form unknown. No budget line, no minister's signature. Hours.

Package C — Elevated reallocation. Act on continuity as if confirmed: advise reserve drawdowns, recast appeal cases, publish the revised outlook prominently. Cost: real-world reserve and appeal decisions staked on one unconfirmed seal, in a domain where the cost of being wrong is counted in the food security of import-exposed populations.

Her senior economist's objection was the professional consensus in one paragraph: *"The disruption case isn't a mood — it's mechanism: payments, shipping, insurance, fertilizer inputs. The carve-outs are untested. And your dated artifact says nothing the pre-war export tables don't. If we reweight and the mechanisms bite next spring, we'll have anchored the world's balance sheet on a tweet."* Touré's draft reply for Thursday: *The mechanisms are real — which is why B moves our analysis toward them, where they bind on the buyer's side, instead of spending another quarter stress-testing the production base of the world's largest exporter in its record year. The reweight isn't faith in a seal; it's consistency with our own facts. And if continuity holds, I want us to have spent the winter thinking about what a strong Russian food sector* does *— because a strangled sector gives nothing away, and a strong one might.* The lock was Thursday; the memo needed a scenario table by then.

EXHIBITS

Exhibit 1 — The consensus mechanism, November 2022. The live fear: payment friction, shipping/insurance barriers, fertilizer-trade disruption reaching Russian food output and world markets — the operational premise behind the grain corridor diplomacy and the fertilizer carve-outs. Facts on the other side of the same table: Russia the #1 wheat exporter entering the war; the record 2022 harvest largely gathered; food/fertilizer deliberately outside the sanctions net.

Exhibit 2 — The sealed input. Public post of November 19, 2022 (tweet 4/23 of a sealed sitting): *"Agriculture and food supply will remain strong in Russia. They might even export these to other countries for profit."* Re-sealed verbatim September 4, 2023, with a resources clause whose outcome was mixed — commodity exports funded the wartime budget at scale *and* 2023 oil-and-gas budget revenues fell 23.9% from the 2022 peak — a counter-receipt printed at full strength on the advisory itself. Documentation and grading at

jyotishintelligence.com/intel/advisory/IA-RU-045. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — What the balance sheet feeds (composite teaching device). Downstream clients: donor appeal sizing; import-dependent reserve authorities (precautionary stocks sized partly on supply tails); alternative-supplier diplomacy. Unit discipline in force: output figures never restated as export claims; bunker weight never as clean weight.

Exhibit 4 — What subsequently happened (public record, for the teaching note). February 2023 (seal + ~90 days): Rosstat confirms the largest grain harvest in Russian history — 157.6 Mt bunker weight (+29.9% y/y), wheat 104.2 Mt (+36%). Season 2023/24: the all-time Russian wheat-export record (~54–55.5 Mt, world #1); 2024/25 holds #1 at 42+ Mt; FAO (June 2025): Russia "will remain the world's largest wheat exporter," ~22% of world markets. Calendar 2023: record $43.5B agricultural export revenue on 66 Mt shipped. July 2023: Moscow pledges 200,000 t of free grain to six African states; first vessel November 7, 2023; final cargo unloaded in Somalia **February 17, 2024 — 455 days after the seal** — the sector the consensus expected strangled operating as an open instrument of statecraft. Fertilizer, in its exact units: output 59.3 Mt (2023), 63 Mt (2024); exports ~31–33 Mt (2022–23), the 40 Mt record only in 2024.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Hold disruption case	B · Quiet reweighting	C · Elevated reallocation
What it adds	consensus safety	a base case matching the desk's own facts + analysis moved to binding constraints + the statecraft page opened	decisiveness
What it forecloses	the ninety-day analytical lead	nothing — it's a memo	reserve safety margins, on one seal
Cost if disruption bites in spring	0 (right with everyone)	a revised memo; analysis already on the binding mechanisms	real harm, publicly owned
Cost / value if continuity holds	rewrite forced by Rosstat, from behind	three seasons anchored right; statecraft class anticipated	right — recklessly

Speed	—	hours to days	committee + client cycle
Authority required	none	unit's own	principals + clients
Reversibility	n/a	full	poor
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/grain-ledger
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — Planning on Continuity

Rails: the momentum objection and the carve-out objection lead and may win — the record itself volunteers both, and "continuity measured sanction design, not foresight" is an admissible verdict; the resources counter-receipt (−23.9% oil-and-gas revenues) stays printed; units are sacred (output ≠ exports; bunker ≠ clean); no welfare outcomes are claimed beyond documented flows; no position on sanctions policy; the free-grain program is analyzed as statecraft mechanics, never celebrated or condemned.

LEARNING OBJECTIVES

(1) Base-case ownership: a balance sheet's scenario weighting is the highest-leverage sentence an analytic unit writes — everything downstream inherits it — and consensus weightings deserve the same interrogation as contrarian ones; (2) mechanism-versus-outcome analysis: the disruption case was built on real mechanisms (payments, insurance, inputs) — the discipline is asking *where those mechanisms actually bind* (buyer-side chokepoints) versus where fear placed them (the production base of a record-harvest exporter); (3) the unthinkable-scenario page: frames don't just weight outcomes, they delete them — a strangled-sector frame cannot *contain* grain-as-statecraft, and the reweight's deepest value was making a scenario class thinkable at all (the Third Shape's sibling, in commodities); (4) counter-receipt reading: the same sealed sitting carried a clause that resolved mixed, printed at full strength — students practice weighting a source whose ledger shows both.

ASSIGNMENT QUESTIONS

1. Write Touré's Thursday scenario-lock memo: the reweighted base case argued from the unit's own facts (the seal cited or not — decide and defend), the redirected analysis plan, the client flag on disruption premiums, and the statecraft scenario page at class level.
2. The momentum and carve-out objections, adjudicated: if the harvest was in the bins and food was outside the sanctions net, what *exactly* did the sealed read add beyond the desk's own tables? Price the honest answer (a dated artifact against the room; the *export-for-profit* extension; agenda-setting) — and state what verdict follows if the answer is "little."
3. Trace the mechanism analysis properly: list the disruption case's channels (payments, shipping, insurance, fertilizer inputs), and for each, determine from Exhibit 4 where it actually bound in 2023–24 — and on whose side of the

transaction. What does the exercise teach about where to spend stress-testing effort under sanctions generally?

4. The free-grain program: reconstruct it as statecraft mechanics — cost to Moscow (~200 kt against 66 Mt shipped), influence purchased across six states, and the analytic predicate required to have staged counter-programming options early. What could a planner holding the continuity read have had ready by July 2023, at class level, without foreknowing any tonnage or recipient?
5. Complete Exhibit 5 and sign it; fill the hindsight row as of February 17, 2024 (the final cargo), honoring the rails — the momentum objection does not disappear because the outcome matched.
6. **The machine staff-officer.** An AI supply-forecasting system, trained on trade flows and sanctions text, would have read the carve-outs and the harvest data and likely weighted continuity *without* any seal — but it would never have opened the statecraft page, because "a sector becomes an instrument" is not a supply variable. What does this split teach about which halves of scenario work machines inherit first, and write the unit's rule for keeping frame-generation (the unthinkable pages) an explicitly human production.

THE BOARD PLAN (75–90 MIN)

Open on Exhibit 3's downstream map (10): the class must feel the balance sheet's leverage — reserves and appeals inherit this sentence. The consensus mechanism (15): build the disruption case at full strength from Exhibit 1; it must be respectable, because it was. The input and the deflations (15): Exhibit 2 with Q2 run immediately — the momentum/carve-out adjudication is the case's honest center. The mechanism-relocation exercise (15): Q3 on the board — where did each channel actually bind? Commit (15): memo + Exhibit 5. Reveal Exhibit 4 (10): Rosstat at +90 days, three seasons of records, the 455-day delivery — units kept exact. Close (10) with Q4's statecraft reconstruction and objective 3 stated plainly: the frame's deleted scenario was the expensive one.

WHAT HAPPENED, AND THE GRADE

The numbers arrived in the continuity column and stayed there for three seasons: history's largest Russian grain harvest confirmed within ninety days of the seal; the all-time wheat-export record the following season; world-number-one held through 2024/25 with FAO projecting more; record export revenue — and then the clause that had sounded most speculative, *"might even export these to other countries for profit,"* exceeded in kind when Moscow delivered 200,000 tonnes of *free* grain to six African states, completing the program 455 days after the seal: the strangled-sector scenario's exact inversion, statecraft where the consensus had priced collapse. The record grades the continuity call a HIT and prints beside it everything this note enforces — the momentum objection, the carve-out objection, and the same sitting's mixed resources clause at full counter-receipt strength. For the composite unit, Package B's value was never the seal's vindication; it was ninety days of analytical lead over

Rosstat, a winter of analysis spent where the constraints actually bound, and one scenario page — grain as instrument — open on the desk when July's pledge made it the only page that mattered.

EPILOGUE

This case teaches the volume's frame-discipline at its largest scale: a whole professional field, for honest mechanistic reasons, weighted a scenario against its own tabled facts — and the weighting deleted the strategic scenario before anyone could price it. The residue for every analytic unit that owns a base case: *once a quarter, argue the weighting from your own tables as if the room's fear didn't exist — and for whatever case wins, write the page your current frame makes unthinkable.* The strange input's role, here as throughout this book, was only to force the exercise onto a dated calendar. The tables were always everyone's to read.

CASE THIRTY-FIVE

Holding the Line: A Counterparty Desk in a Short-Seller Storm

A teaching case in credit risk — the asymmetry between a hold and an unwind, priced at the panic's worst prints

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the sealed forecast in this case is graded on that record's open ledger as a **NEAR, not a hit**: three of its four legs landed, the fourth was cut by an exogenous legal shock, and a live upgrade trigger has not fired as of this writing — nothing in this case depends on it firing. The record's disclosures are inherited in full: its publisher states it holds no position in, and has received no compensation relating to, any Adani security; the ₹3,456 level is a grading threshold, not a price target; **nothing in this case is investment advice**; and a sealed recovery forecast is **never a license to relax due-diligence obligations** — the boundary the case itself polices. The strongest objection travels up front: survival of a systemically important, politically connected group may have been the smart-money base case regardless of any forecast, and the statement carried no numeric probability. Every price, purchase, ruling, indictment, and settlement is real and cited. The **protagonist and his bank are composite characters** — "Farhan Qureshi" and "Straits Commercial Bank" are not real; their book and committee are teaching devices, and no outside-world fact discussed is invented. The case takes **no position** on any allegation against any party — the U.S. criminal charges were denied and later moved to be dropped — and none on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

THURSDAY, FEBRUARY 9, 2023 — 08:00, SINGAPORE

Farhan Qureshi ran counterparty credit for Straits Commercial Bank's South Asia book, and for two weeks his mornings had begun with the same name. The Hindenburg Research short report on the Adani Group — fraud allegations, denied by the group — had erased more than **$150 billion** of

market value in fourteen trading days; the flagship had withdrawn a fully subscribed **$2.4 billion** share sale; and the consensus across every desk Qureshi respected had hardened into two words: *terminally impaired.*

His exposure surface was the standard one, and every element of it demanded a posture this week: **collateral eligibility** (group paper securing trade-finance facilities — keep accepting at a wider haircut, or strike it from the list?); **facility posture** (two facilities rolling in March — renew, tighten, or decline?); **limits** (hold, trim, or freeze new business?); and **margin sequencing** (the crash itself was forcing calls — enforce on schedule, or accelerate?). All of it sat inside his committee's delegated authority. None of it required the bank to decide whether the allegations were true — only what its book would do while the world found out.

The item his analyst flagged — *"contrarian, dated, and public, for whatever that's worth"* — was a timestamped post of **February 8**, from a source with a sealed and publicly graded record, method disclosed as Vedic astrology. Four separable claims: the group *"has a good future"*; the stock would *"recover/trend upwards on average starting March 21"*; it would *"continue to fluctuate until April 2024"*; and *"from May 2024, it will begin a new journey of even greater growth."*[1] Survival, a dated turn, a bounded volatility period, a dated growth phase — each falsifiable against the public tape, posted where being wrong would be permanent.

Qureshi's file note was a credit officer's, not a believer's. The consensus was priced at panic prints; the statement was the only dated counterweight in his intake; and the decision menu in front of him had a structural asymmetry that mattered more than either: **a hold was close to free** — no transaction, no crystallized loss, reversible by memo — while **an unwind at the bottom** converted mark-to-market drawdown into realized losses at the worst available prints and spent franchise capital with a systemically important client that could take years to repurchase. The question the sealed post posed was not "is this true?" It was: *at this asymmetry, what does it cost me to let a dated, falsifiable, contrarian artifact stiffen a posture my own arithmetic already favors — and where is the boundary beyond which that becomes negligence?*

THE THREE POSTURES

Package A — Drift with the consensus. No decision is a decision: let the panic set the posture — de-list the collateral, decline the rolls, freeze the limits. Cost: realized losses at the lows, a de-listing signal amplified in a fragile

market, and a franchise exit executed at maximum spread. Defensible to any committee in that week — which was exactly what made it expensive.

Package B — The quiet hold. Retain collateral eligibility at a **punitive haircut**; roll the March facilities with **tightened covenants**; **hold** limits while freezing only new business; enforce margins **on schedule, without acceleration** — prudence without pro-cyclical panic. Every element inside existing authority, executable in days, reversible if the credit picture deteriorated on *credit* facts. Documented in a dated committee memo whose stated logic was the asymmetry itself, with the sealed post logged — honestly — as one input among many, and the boundary line written in: *this posture holds unless due diligence surfaces credit-relevant confirmation of the allegations; no forecast overrides that.*

Package C — The conviction posture. Treat the recovery forecast as information: widen limits, add exposure at distressed levels, position the bank for the rebound. Cost: the category error the whole case exists to name — a counterparty desk trading its client's crisis on an astrologer's calendar, due diligence subordinated to a forecast, and every fiduciary and compliance boundary crossed at once.

The chief risk officer's challenge, when Qureshi previewed the memo: *"Strip out the stargazing and tell me why B isn't just wishful thinking with a haircut. The whole street says impaired."* Qureshi's answer, as drafted: *B stands on three legs that have nothing to do with the post: the asymmetry (holds are cheap, bottom-unwinds are not), the client's systemic footprint (politically connected, infrastructure-critical — survival is arguably the base case, which cuts against the street's certainty, not for the post's), and reversibility (every B element can become A in a week if credit facts arrive). The post contributes exactly one thing: it is dated and falsifiable where the street's certainty is neither — and if its March date passes with no turn, that failure is* also *information, and my memo says what we do with it.* The committee sat Monday.

EXHIBITS

Exhibit 1 — The episode, day fourteen. Hindenburg report (Jan 24): fraud allegations, denied; >$150B market value erased; the $2.4B share sale withdrawn; flagship stock approaching what would be its ₹1,017 low; consensus: terminally impaired.

Exhibit 2 — The sealed input. Public post of February 8, 2023 (platform-timestamped): the four legs quoted verbatim in the case text. Advisory documentation and grading at jyotishintelligence.com/intel/advisory/IA-MKT-001 — grade **NEAR** (three legs landed; the fourth cut by the November 2024 indictment); live upgrade trigger ₹3,456, **unfired**; publisher's no-position/no-compensation disclosure; the threshold is a grading level, not a price target. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The decision surface (composite teaching device). Collateral: group paper at standard haircut, eligible. Facilities: two rolling in March. Limits: full; new business active. Margins: calls in progress. Authority: all four levers inside delegated committee authority; no board action required for B; C requires appetite the mandate does not contain.

Exhibit 4 — What subsequently happened (public record, for the teaching note). Late Feb 2023: flagship bottoms near ₹1,017. Early March 2023: GQG Partners buys ~$2B across four group companies — the turn arrives, **ahead of the named March 21**. Fluctuation through the SEBI probe; January 2024: the Supreme Court declines to transfer the investigation. **May 24, 2024:** Adani Enterprises erases the last of its Hindenburg losses at ₹3,456 — in the month named fifteen months earlier. **November 2024:** U.S. federal indictment on alleged bribery charges (denied); dollar-bond access restricted; the stock falls back below the May high — **the fourth leg breaks.** By 2026: DOJ moves to drop the criminal charges; SEC (~$18M) and Treasury/OFAC (~$275M) matters settled; shares recovering, the May-2024 peak not decisively reclaimed. Grade: NEAR, upgrade trigger unfired.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Drift with the consensus	B · Quiet hold	C · Conviction posture
What it adds	committee safety in week two	the cheap side of the asymmetry, documented	upside exposure
What it forecloses	re-entry for quarters/years	nothing — reversible by memo	the mandate, the boundary, possibly careers
Cost if the group collapses	vindicated, at the cost already paid	losses bounded by haircuts and covenants — and a memo that aged badly	ruin, deserved

Cost / value through the actual path	realized losses at ₹1,017-era prints; franchise exit	collateral marked ~3× by May 2024; client retained; posture survived even the indictment leg	profits the mandate never authorized
Speed	days	days	days
Authority required	committee (by drift)	committee	none that exists
Reversibility	poor	full	n/a
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading — the record itself says NEAR)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/counterparty-window
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — Holding the Line

Rails: the grade is **NEAR and stays NEAR** — three legs, a broken fourth, an unfired trigger; no telling rounds it up; the smart-money objection is conceded and load-bearing (survival of a systemically important group may have been the sophisticated base case — the post's contribution must be argued *net* of it); the allegations are neither adjudicated nor dismissed by this case — the desk's boundary line ("no forecast overrides due diligence") is the pedagogy; the no-position, no-advice, threshold-not-target disclosures are honored in every discussion; nothing herein is investment advice.

LEARNING OBJECTIVES

(1) The hold/unwind asymmetry as the master structure of crisis credit decisions: cheap, reversible postures versus expensive, irreversible ones — and why panic consensus systematically pushes committees toward the expensive side at the worst prints; (2) the boundary between input and license: a contrarian forecast may stiffen a posture the desk's own arithmetic favors, and must never substitute for due diligence — students draw the line in writing, in the memo itself; (3) multi-leg grading discipline: consuming a four-leg call whose legs resolve separately (turn early, window held, peak on the named month, growth leg severed by an exogenous shock) — and what a NEAR with a live trigger teaches about grades as ongoing objects; (4) the falsifiability dividend: the post's one uncontestable contribution — a dated claim whose *failure* would also have been usable — versus the street's undated certainty, which could never be wrong on schedule.

ASSIGNMENT QUESTIONS

1. Write Qureshi's Monday committee memo: the four levers set, the asymmetry argued from the bank's own numbers, the sealed post logged honestly among inputs, the boundary clause verbatim, and the March-date failure protocol (what the desk does if the named turn doesn't come).
2. The smart-money objection, priced: construct the survival base case from February 2023 facts alone (systemic footprint, political economy, asset quality). What does the sealed post add *net* of that case — and is "a dated artifact against an undated consensus" worth anything when the sophisticated minority already agreed? Argue both ways.
3. The four legs, graded separately: turn (early — does early count?), volatility window (held), May 2024 recovery (exact month, fifteen months out), growth leg (severed by an indictment no market model contained). What does this decomposition

teach about how exogenous legal shocks should interact with forecast grading — and would you have designed the record's NEAR-plus-trigger treatment differently?

4. Package C is the case's cautionary tale, but its boundary is subtler than it looks: a *treasury* desk buying distressed bonds on its own analysis is legitimate; a *counterparty* desk widening limits on a forecast is not. Locate the exact difference (mandate, information role, whose risk), and write the compliance rule that permits the first while barring the second.
5. Complete Exhibit 5 and sign it; fill the hindsight row as of December 2024 (post-indictment) and again as of mid-2026 (charges dropped, settlements done, trigger unfired). Does B survive both readings? That dual test is the point.
6. **The machine staff-officer.** An AI credit-surveillance system ingests news, prices, and sealed forecasts, and in week two of a short-seller episode recommends postures. Trained on episodes where allegations proved true, it recommends A; trained on survivals, B; and its confidence in either is an artifact of its corpus. What does this case teach about base-rate poverty in tail episodes (every short-seller crisis is nearly unique), and which memo elements — the asymmetry argument, the boundary clause, the failure protocol — must remain human-authored precisely because no corpus prices them? Write the model-governance paragraph.

THE BOARD PLAN (90 MIN)

Open on Exhibit 3 (15): the four levers, week two, and the committee pressure — students must feel how *defensible* A is before B can be interesting. The asymmetry (15): objective 1 built from first principles, no forecast yet. The input (15): Exhibit 2 with the NEAR grade and disclosures read aloud; run Q2's net-contribution fight. The boundary (15): Q4's treasury-versus-counterparty distinction — the case's sharpest edge. Commit (15): memo + Exhibit 5. Reveal Exhibit 4 (10): GQG early, May 2024 exact, the indictment breaking the fourth leg, the 2026 resolutions — the untidy full path. Close (5) on Q3's grading debate and the record's own verdict: NEAR, trigger live, nothing rounded up.

WHAT HAPPENED, AND THE GRADE

The path was neither the street's nor the seal's, which is why it teaches. The turn came in early March 2023 — GQG's ~$2 billion arriving ahead of the named March 21 — the stock fluctuated through the SEBI probe and a Supreme Court ruling exactly as the bounded-volatility leg described, and on **May 24, 2024**, fifteen months after the seal, Adani Enterprises erased the last of its Hindenburg losses in precisely the month the post had named. Then November 2024 severed the fourth leg: a U.S. indictment — charges denied, later moved to be dropped, with SEC and OFAC matters settled by 2026 — cut the stock below its May high, where its recovery has since hovered without decisively reclaiming the peak. The record grades the call **NEAR and leaves it there**: three legs landed, one broken by an exogenous shock, the ₹3,456 upgrade trigger published, live, and unfired — no hit claimed, the smart-money objection printed in

its own file. For the composite desk, Package B's ledger was the asymmetry realized: collateral held at punitive haircuts was marked roughly three times higher by May 2024, the client relationship survived into the recovery, and the posture — because it was built on the bank's own arithmetic with the forecast as a logged input, never a license — survived even the leg the forecast lost.

EPILOGUE

This case closes the volume's markets thread — the Tariff Shock's committee, the City Hall Window's scorer, the gilt desk's selectorate, and now the counterparty officer — and its residue binds them all: *in a panic, the expensive decision is usually the defensible one, and the cheap decision needs a document.* The sealed post's deepest contribution was never its dates. It was that it could be wrong on a schedule — and a desk that files its own postures with that same property, dated logic, named falsifiers, failure protocols, has learned from the record the only thing this book has ever really taught.

CASE THIRTY-SIX

The Four-Day Fuse: Reserve Traction Against a Dated Window

A teaching case in operational logistics — converting an undated expectation into a shift plan, and buying the first rehearsal before the first wave

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the sealed call in this case is graded **PARTIAL** on that record: its timing fuse and friction clauses landed; its strongest clause — *successful* interdiction — **failed**, refuted by the target railway's own bulletins and by the adversary's own leader, and the record keeps the failure at full weight. Two pre-seal discounts also travel from the first page, stated in the advisory itself: a same-scale strike wave hit **four days before the seal**, elevating the base rate for "strikes within days"; and the railway's diesel-fallback doctrine was **already public** before the seal — so the residual value claimed is only the increment a dated fuse adds over an undated expectation. Every delay figure, recovery time, and bulletin is real, cited from the railway's own operating log. The **protagonist is a composite character** — "Petro Lysenko" is not a real person; no actual railway's internal deliberations are depicted, and no outside-world fact discussed is invented. The real railway's performance in these events was, by any standard, extraordinary, and this case grades nothing about it. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

SATURDAY, NOVEMBER 19, 2022 — 20:00, A NATIONAL RAILWAY OPERATIONS CONTROL CENTRE

Petro Lysenko owned the night shift in the operations control centre of a railway at war — traction assignment, reroute execution, service continuity across a network that was, functionally, the country's supply artery. His professional world had two facts nailed to it. **The network was electrified**, nationwide, which meant the adversary's missile campaign against the power grid was, from his chair, an attack on traction power. **The fallback was diesel** — reserve locomotives, held at depots, swapped in when the contact wire died

— and the doctrine was not merely internal knowledge: the railway's own CEO had posted it publicly two weeks earlier (*"when they shell our energy infrastructure, we switch to diesel locos and keep running"*), and an international wire profile of the railway's wartime performance had run three days ago.

Four days earlier, on November 15, a mass wave had hit the grid — so nobody in the building needed convincing that more strikes were coming. *More strikes are coming* was the ambient condition of the job. What the ambient condition could not tell him was **when**, and *when* was the variable his shift actually ran on: which depots held the reserve diesels tonight, which junctions had fuel, which diesel-qualified crews sat on which rosters.

The item a colleague showed him — with a shrug — was tweet 9 of a public 23-tweet thread posted that day, from a source with a sealed, graded record, method disclosed as Vedic astrology: *"In a few days from now, Russia will be able to successfully obstruct military and other supplies to Ukraine. Supplies can no longer easily reach Ukraine without obstruction from Russia."*[1] Lysenko read it as a professional reads any warning: in clauses. A **timing fuse** — "in a few days." A **friction claim** — supplies no longer moving easily. And a **strong interdiction claim** — "successfully obstruct" — which his own railway's public record already argued against: the fallback existed, was drilled, and was published. The strong clause he discounted on the spot. The fuse was the only content the ambient expectation lacked — and a fuse, unlike an expectation, was **rosterable**. The difference between "more waves are coming" and "the window opens within days" was the difference between a standing posture and a shift plan.

THE THREE POSTURES

Package A — Standing posture. The doctrine exists; the depots are stocked; deploy reactively, wave by wave, as the doctrine says. Cost: zero. Exposure: the first wave's response paid for its own rehearsal — traction swaps, fuel moves, and crew calls executed for the first time *during* the event, with the delay envelope carrying the learning cost.

Package B — Pre-stage against the fuse. Three actions inside duty-officer/operations-director authority, executable within the four-day window, priced in shunting movements and overtime. *One:* move reserve diesel traction from depots toward the network's **known continuity chokepoints** — not any location a strike would later reveal (the seal named no line, no corridor), but the ordinary junctions a duty officer already prioritizes. *Two:* pre-stage **fuel** at those junctions and roster **diesel-qualified crews** onto the

shifts covering the window — lead times of days, exactly what the fuse supplied. *Three:* pre-draft **reroute plans** and open a standing coordination line with grid dispatch on general voltage risk — planning hours, no foreknowledge of any specific line required.

Package C — Network-wide maximal posture. Hold trains, pre-cancel services, demand grid-side priorities — treat the fuse as certainty and the strong clause as credible. Cost: self-inflicted service disruption ahead of any strike, doing a fraction of the adversary's work for him.

The operations director's question, when Lysenko called it in, was the shift-economy one: *"We've been at this for months. The doctrine works. What does four days of shunting buy me that Wednesday's wave didn't already teach us?"* Lysenko's answer, as he logged it: *The doctrine tells us what to do; it doesn't tell us when to have already done it. Every wave so far, the fallback's first hours were spent deploying — that's time in the delay column, and it's the one variable we control in advance. The post is worth nothing on 'whether.' On 'when,' it's the only dated artifact anyone's offering — and everything it makes me do is work we'd do anyway, four days later, under fire.* The shift plan for the window was his to write that night.

EXHIBITS

Exhibit 1 — The ambient condition, November 19, 2022. November 15: a same-scale mass wave against the grid — four days *before* the seal; base rate for further strikes: elevated, universally assumed. The railway's fallback doctrine public: the CEO's diesel-locos post (November 6); the wire-service profile (November 16; 85% of trains on schedule the prior month). The live analytic debate: not *whether* strikes, but whether the campaign could collapse grid and logistics over the winter.

Exhibit 2 — The sealed input. Public post of November 19, 2022 (tweet 9/23, ID 1594080616085348352), quoted verbatim in the case text — three separable clauses: fuse, friction, strong interdiction. Advisory documentation and grading at jyotishintelligence.com/intel/advisory/IA-RU-034: **PARTIAL** — fuse and friction landed; "successfully obstruct" failed, kept at full weight; both pre-seal discounts stated in the record's own file. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The duty officer's levers (composite teaching device). Inside shift/operations-director authority: reserve-traction positioning (shunting + scheduling); fuel logistics to junctions; crew rostering (overtime); reroute pre-drafting; grid-dispatch coordination line. Outside it: anything touching service cancellation or grid priorities.

Exhibit 4 — What subsequently happened (public record — the target railway's own bulletins). November 23 — four days after the seal: ~70 cruise missiles and 5 drones against critical infrastructure; grid frequency collapse trips **every operating nuclear plant** offline (Rivne, South Ukraine, Khmelnytskyi); large parts of the rail network de-energized; reserve diesels deployed; 13 long-distance services rerouted; next morning's bulletin — *"95 of our trains keep moving despite everything… 81 of them are running more than an hour late"* — worst case +12h51. The waves repeat: December 5 (worst +4h27, voltage back by morning), December 16 (four oblasts de-energized, worst +9h24), December 29 (69 missiles, 23 drones — and by afternoon just **6 of 78 services** more than 30 minutes late, recovered same-day). Mid-campaign, the railway: *"not a single long-distance train has been cancelled; reserve diesel locomotives were brought into service instantly."* ISW, day one: "a misguided attempt to degrade the Ukrainian will to fight"; CSIS, May 2023: the missile war "severely underperformed expectations." The adversary's own leader, March 2023, on why deliveries couldn't be intercepted: the equipment moves "covertly, at night… what we manage to reach, we destroy. That goal certainly exists" — pressure conceded, interdiction not achieved.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Standing posture	B · Pre-stage against the fuse	C · Maximal posture
What it adds	nothing — doctrine suffices	the first wave's rehearsal bought in advance	certainty theater
What it forecloses	the deployment hours in wave one's delay column	shunting, fuel moves, overtime	service, preemptively — the adversary's goal, self-delivered
Cost if the window passes quietly	0	assets repositioned, a quiet week's overtime	cancelled trains for nothing
Cost / value when the wave lands	fallback deploys under fire, learning on the clock	traction, fuel, crews already at the junctions	protected — at the network's own expense

Speed	—	inside the 96-hour fuse	days, with authority fights
Authority required	none	duty officer / operations director	board + government
Reversibility	n/a	full (shunt them back)	poor
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading — the record says PARTIAL)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/reserve-traction
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Four-Day Fuse

Rails: the PARTIAL is the spine — the strong clause **failed**, refuted by the target's own bulletins and the adversary's own words, and no telling rounds the call up; both pre-seal discounts (elevated base rate; public fallback doctrine) are conceded at full strength — the residual is *a dated fuse versus an undated expectation*, and nothing more; the real railway's performance is treated with the respect it earned and graded by no one; the composite officer's decisions are the only decision surface; wartime human costs are acknowledged and never converted into teaching arithmetic.

LEARNING OBJECTIVES

(1) The fuse increment: the precise, narrow value of *when* against an already-believed *whether* — a dated window converts standing doctrine into a shift plan (rosters, fuel, positioning), and the increment is measured in the first wave's deployment hours, nothing grander; (2) clause-graded consumption in real time: the officer discounted the strong clause *on the spot*, using his own institution's public record as the counter-evidence — the volume's decomposition discipline exercised at shift speed; (3) rehearsal economics under fire: the railway's own wave-by-wave trajectory (81-of-95 late → same-day recovery by wave four) is the cleanest documented evidence in this book that response speed is a *learnable* variable — and a dated warning is a way of buying rehearsal number one before the event rather than during it; (4) grading asymmetric outcomes: a call whose fuse landed exactly, whose friction is documented from the target's own log, and whose headline verb failed — the PARTIAL as the honest composite, and the discipline of letting the failed clause cap the claim.

ASSIGNMENT QUESTIONS

1. Write Lysenko's shift plan for the window: traction moves (to which class of junctions, on what logic), fuel and crew rosters, the reroute drafts, the grid-dispatch line — every item bounded to what the sealed clauses support, with the strong clause's discount logged in the plan itself.
2. Price the fuse increment honestly: using Exhibit 4's wave trajectory, estimate what fraction of the first wave's delay envelope was deployment time rather than damage time — and therefore what pre-staging could and could not have compressed. Where does the "might have" ceiling sit, and why can it go no higher?
3. The strong clause failed against exactly the resilience that was publicly visible before the seal. What does this teach about forecasts that name an *outcome* ("successfully obstruct") versus forecasts that name a *window* — and which clause class should operational desks systematically weight?

4. The operations director's question — "what does four days of shunting buy that Wednesday's wave didn't teach?" — has a real answer and a real limit. Give both: what the November 15 experience supplied (doctrine validation), what it couldn't supply (the next window's date), and what neither supplied (the strong clause's truth value).
5. Complete Exhibit 5 and sign it; fill the hindsight row as of December 30 (wave four's same-day recovery), crediting the trajectory to the railway's learning — and stating honestly what share, if any, a pre-staged first wave could have claimed.
6. **The machine staff-officer.** An AI operations system now watches launch-platform activity, fuel logistics, and adversary signals, and issues probabilistic strike-window estimates continuously — fuses on tap, no astrologer required. The fuse increment becomes permanent. What new failure mode arrives with it (alert fatigue? pre-staging as the constant state, which is just a bigger standing posture)? Write the duty-room rule that keeps machine-fused windows converting into *distinct shift plans* rather than ambient anxiety.

THE BOARD PLAN (75 MIN)

Open on Exhibit 1 and the ambient condition (10): everyone in the room already believes strikes are coming — the class must feel how little "more strikes" is worth, so the fuse can be worth something. Clause surgery at shift speed (15): Exhibit 2 decomposed; the strong clause discounted with the railway's own public doctrine — objective 2 done live. The fuse-versus-expectation distinction (15): objective 1 built on the board — what specifically becomes possible when "soon" becomes "within days"; list the rosterable items. Commit (15): the shift plan + Exhibit 5. Reveal Exhibit 4 (15): the wave, the nuclear trips, 81-of-95, the wave-by-wave learning curve to same-day recovery — read from the bulletins; then the interdiction verdict from ISW, CSIS, and the adversary's own mouth. Close (5) with Q3's clause-class rule: windows over outcomes, always, at operational desks.

WHAT HAPPENED, AND THE GRADE

The wave came on the fuse's schedule: November 23, four days after the seal — seventy cruise missiles, the grid staggered, every operating nuclear plant tripped, great sections of an electrified railway suddenly without traction power. The friction was real and documented from the target's own log: eighty-one of ninety-five long-distance trains more than an hour late the next morning, the worst nearly thirteen hours. And the strong clause died against the fallback: reserve diesels in service "instantly," not one long-distance train cancelled, every wave cleared in 24–48 hours, and by the fourth wave — December 29, sixty-nine missiles — the railway was recovering same-day, a learning curve written in its own bulletins. The record grades the call **PARTIAL** and holds it there: fuse exact, friction documented, interdiction refuted by the railway's log, the analysts' verdicts, and the adversary's own concession that what moves at night mostly gets through. The two discounts stay printed beside it. For the composite duty officer, Package B's claim was never victory — the railway's

people own that entirely. It was four days of shunting, fueling, and rostering that would have moved rehearsal number one ahead of wave number one, and a delay column, in the first hard morning, shorter by exactly the hours that preparation is worth.

EPILOGUE

This case is the volume's purest specimen of the smallest unit of foresight value: not a prediction believed, but **a date attached to an expectation everyone already held** — worth precisely one rehearsal, bought early. The railway's bulletins, read as a series, are the proof that rehearsals compound; the failed strong clause is the proof that outcomes should never have been the clause anyone consumed. For every operational desk that lives on doctrine and fallbacks, the residue is one line for the control-room wall: *doctrine answers "what"; only a date answers "when to have already finished."*

CASE THIRTY-SEVEN

A Named Actor, Sixteen Days Early: The Ministry SOC

A teaching case in cyber defense — what an actor-and-window claim is worth to detection engineering, and what it can never claim to prevent

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the strongest objection travels first: the sealed line is **one item inside a thirty-one-item single thread** naming many actors, mechanisms, and domains in one sitting, and a hostile reader is entitled to ask whether a long, multi-domain thread simply covers enough ground that some line resolves true. That coverage objection is conceded, not argued away. The specific delivery mechanism later documented (EnvyScout, ISO/IMG lures) was **never in the seal** and is treated strictly as outcome; **no causal chain is claimed** from any ministry's 2023 posture to the same actor's 2024 intrusions at Microsoft and HPE; and the honest ceiling of every counterfactual clause is **raised attacker cost, never prevented use**. Every attribution, disclosure, and advisory is real and cited. The **protagonist and her ministry are composite characters** — "Katrin Vogel" and her SOC are teaching devices; no actual government's network operations are depicted, and no outside-world fact discussed is invented. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

MONDAY, FEBRUARY 27, 2023 — 08:30, AN EU FOREIGN MINISTRY

Katrin Vogel ran the security-operations centre of a European foreign ministry whose diplomatic traffic ran thick with Ukraine coordination — exactly the target profile a foreign-intelligence service builds collection against. Her SOC's posture, like the industry's attention, had rotated over the winter: the loud clusters were the kinetic, GRU-attributed ones — Sandworm, APT28 — wipers and infrastructure attacks tracking the war's rhythm. The quieter SVR-attributed espionage cluster — **Nobelium** in Microsoft's naming; APT29, Cozy Bear to the rest of the industry — had drifted, for months, out of

the trade press and, if she was honest, some notches down her own detection-engineering queue. Tradecraft rules aged; staff advisories referenced last year's lures.

The item her threat-intel analyst brought to Monday triage was, by every conventional standard, inadmissible: a public X thread of **February 26** — thirty-one items, many actors and domains in one sitting, from a source whose record was sealed and publicly graded, method disclosed as Vedic astrology. Item T23 named her world precisely: *"cyber-espionage and hacking attacks on adversaries via the Nobelium group"* assessed likely, inside a window running to **April 30**.[1] A specific Microsoft cluster designation. An operation class. A nine-week window. No target, no lure, no technique.

Vogel logged the two honest readings side by side. *Against:* a thirty-one-line thread is a shotgun — name enough actors and something attributes eventually; the coverage objection wrote itself, and she wrote it. *For:* the line was **divergent** — the room's attention was on the GRU clusters, and nobody in her feeds was flagging Nobelium this quarter — and it was **actionable at exactly her desk's grain**: detection engineering doesn't need a lure named; it needs an actor named, because actors carry documented, public, *already-known* tradecraft. The question was the volume's usual one, wearing a SOC badge: what did the line make cheap — and was any of it work her queue didn't already owe?

THE THREE POSTURES

Package A — Monitor. Note the thread; let the queue stand. Cost: zero. Exposure: the posture stays generic until some vendor's public attribution does the naming — after a campaign, not before it.

Package B — Quiet readiness. Three actions inside her standing authority, same-week, near-zero incremental cost. *One:* **refresh and re-tune detection** against Nobelium/APT29's already-public, pre-2023 tradecraft — the spear-phishing, credential-harvesting, and cloud-identity techniques documented across prior CISA and industry reporting on this exact actor. Not new tooling; re-tuning against a known repertoire. *Two:* a **staff advisory** naming the actor cluster and the window — general vigilance, no specific lure theme, since none was named. *Three:* **raised scrutiny on ministry mail traffic** for the window's duration. All of it bounded to the sealed words: actor, class, window.

Package C — Elevated posture. Treat the line as warning of a specific campaign: emergency change windows, cross-agency alerts, leadership briefings citing the source. Cost: the SOC's credibility spent on a thirty-one-line thread, and an alert-fatigue tax on the whole ministry for nine weeks.

Her deputy's objection was the profession's, sharpened: *"If the queue owes APT29 rule-refreshes anyway, do them because the queue owes them — not because line twenty-three of an astrologer's thread said so. And if the queue doesn't owe them, one line shouldn't jump it."* Vogel's answer, drafted for the triage log: *That's the right frame and the wrong conclusion. Queues are priority fights, and 'owed anyway' work loses them without a date attached — that's why our Nobelium rules aged in the first place. The line gives me an actor, a class, and a nine-week window: enough to justify a bounded re-prioritization I can defend in one sentence without mentioning the source — 'SVR-attributed espionage against Ukraine-supporting ministries is our top target profile; our detections against its known tradecraft are stale; fixing that is this week's work.' Every word of that is true independently. The thread's contribution is that it's dated — and if April 30 passes with nothing, the window closes and the queue reverts, which is more than 'someday' ever does.* The triage decision was hers to log that morning.

EXHIBITS

Exhibit 1 — The attention rotation, winter 2022–23. Trade coverage and defensive priority rotated toward kinetic GRU-attributed clusters (Sandworm/APT28); the SVR-attributed Nobelium/APT29 cluster out of the headlines for months; the composite SOC's actor-specific rules aging accordingly.

Exhibit 2 — The sealed input. Public thread of February 26, 2023 — 31 items; item T23 quoted verbatim in the case text: Nobelium named, espionage class, window to April 30. Advisory documentation and grading at jyotishintelligence.com/intel/advisory/IA-RU-002 — graded HIT on all rubric axes; the coverage objection (one line in a 31-item sitting) conceded in the record's own file; misses elsewhere in the same Russia file retained at full weight. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The SOC's authority map (composite teaching device). Inside the SOC lead's discretion: detection-rule refreshes, staff advisories, mail-traffic scrutiny levels, bounded re-prioritization of the engineering queue. Requiring higher sign-off: cross-agency alerts, emergency changes, anything citing sources in leadership channels.

Exhibit 4 — What subsequently happened (public record, for the teaching note). March 14, 2023 — sixteen days after the seal: BlackBerry Threat Research attributes a live phishing campaign against EU governments assisting Ukraine to APT29/Nobelium — EnvyScout/ROOTSAW malware, diplomatic-invitation lures impersonating the Polish ambassador's Washington travel — inside the sealed window. The cluster's later record, **causally unconnected to any ministry's 2023 posture and claimed for nothing here:** January 19, 2024 — Microsoft discloses a Midnight Blizzard breach of its own corporate email, reaching senior leadership; January 24, 2024 — HPE's 8-K discloses an analogous intrusion; February 29, 2024 — CISA/FBI/NSA joint advisory AA24-057A formally attributes the cluster, across all aliases, to Russia's SVR. Four independent organs converging on the actor named in the February seal. October 2024: the actor returns with an RDP-file spear-phishing wave — the reusable-capability point.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · Quiet readiness	C · Elevated posture
What it adds	nothing	tuned detections vs known tradecraft + a named-actor advisory + watched mail, all window-bounded	ministry-wide alarm
What it forecloses	the pre-attribution head start	a week of queue re-prioritization	credibility; nine weeks of alert fatigue
Cost if the window passes empty	0	refreshed rules the queue owed anyway; a reverted priority	fatigue spent on line 23 of 31
Cost / value if the actor operates in-window	generic posture meets a tuned adversary	raised attacker cost — the honest ceiling; earlier-detection odds up	braced loudly, no better tuned
Speed	—	same week, one duty cycle	cross-agency clocks
Authority required	none	SOC lead's own	CISO + leadership
Reversibility	n/a	full (window closes, queue reverts)	poor
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/ministry-soc
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — A Named Actor, Sixteen Days Early

Rails: the coverage objection leads and is conceded — one line in a 31-item thread, and the record says so itself; the EnvyScout mechanism is **outcome-only** — a February reader could tune against *known* APT29 tradecraft, never against the specific lure; the ceiling is **raised attacker cost, never prevention** — no claim that any campaign would have been stopped; **no causal chain** to the 2024 Microsoft/HPE breaches — they demonstrate the actor's persistence, not any ministry's failure; the deflationary reading ("the queue owed this work anyway") is admissible and may win.

LEARNING OBJECTIVES

(1) Actor-grain actionability: cyber defense is the volume's cleanest illustration that different claim grains serve different desks — a *technique* claim serves incident response, but an *actor + window* claim serves detection engineering exactly, because actors carry public, documented, tunable repertoires; (2) the queue-date mechanism: "owed anyway" work loses priority fights until a date attaches — the sealed window functions as a bounded re-prioritization warrant, defensible without the source, reverting when the window closes (the Festival Window's audit logic and the Artemis check's pedigree-free justification, at a SOC); (3) coverage-effect discipline: how to weight one resolving line from a many-line sitting — the objection, the record's own concession of it, and what (divergence from the rotated consensus; the absence of competing public namings) partially answers it; (4) cost-raising as the honest metric of defensive value: students learn to write counterfactuals whose ceiling is the adversary's raised cost of operating, and to refuse every prevention narrative the outcome tempts them toward.

ASSIGNMENT QUESTIONS

1. Write Vogel's triage-log entry and the one-sentence queue justification that never mentions the source — then the internal footnote that honestly records what prompted the timing. Defend keeping both.
2. The coverage objection, quantified as far as honesty allows: a 31-item thread names many actors and domains. What evidence would distinguish shotgun coverage from signal (per-line priors? the record's own grading of the sitting's other lines? competing public namings?), and what does Exhibit 2 actually supply toward that test?

3. Enumerate what "refresh detection against known APT29 tradecraft" concretely meant in February 2023 — technique families from public reporting, no EnvyScout allowed — and mark where each would or wouldn't have intersected the campaign BlackBerry documented. Where does raised-cost value survive even the misses?
4. The 2024 chapter: Microsoft, HPE, the joint advisory, the October return. None of it is causally connected to any 2023 ministry posture — so what *is* its legitimate use in Vogel's story? (Expected: actor persistence justifies reusable capability; the window's work compounds.) Write the boundary sentence that keeps the use honest.
5. Complete Exhibit 5 and sign it; fill the hindsight row as of May 1, 2023 (window closed), in the world as it happened *and* in a world where BlackBerry's attribution never came.
6. **The machine staff-officer.** An AI detection-engineering assistant continuously tunes rules against every tracked actor's evolving public tradecraft — the queue never ages, and the actor-grain advantage becomes ambient. What of this case survives (the window's re-prioritization logic? the coverage-objection discipline for the *AI's own* many-line threat feeds)? Write the SOC's rule for how machine-maintained detection consumes dated, single-source actor claims — including the revert clause.

THE BOARD PLAN (75 MIN)

Open on Exhibit 1's rotation (10): attention as a defensive resource that drifts — students should see the aging queue as the normal condition, not negligence. The input with the objection attached (15): Exhibit 2, coverage problem stated first; run Q2's distinguishing-evidence design. The actor-grain insight (15): objective 1 on the board — map claim grains (actor / technique / target / date) to desk functions; this is the case's transferable core. Build Package B and the pedigree-free justification (15): Q1 drafted live. Commit (10): Exhibit 5. Reveal Exhibit 4 (15): BlackBerry at sixteen days, then the 2024 convergence — with Q4's boundary sentence enforced as it's read. Close (5) on the honest ceiling: raised cost, never prevention — the sentence every cyber after-action should have to carry.

WHAT HAPPENED, AND THE GRADE

Sixteen days after the seal, BlackBerry attributed a live campaign against EU governments assisting Ukraine to exactly the cluster the February line had named — diplomatic lures, EnvyScout delivery, none of it nameable in advance and none of it claimed. The actor's subsequent record wrote the persistence argument no defender needed persuading of: Microsoft's own corporate email in January 2024, HPE days later, a three-agency joint attribution to the SVR that February, a fresh spear-phishing wave that October — four independent organs converging on the designation the seal had used, across thirteen months, with no causal thread to any ministry's posture and none asserted. The record grades T23 a **HIT** on all rubric axes, concedes the coverage

objection in its own file, and keeps its Russia-file misses beside it. For the composite SOC, Package B's ledger was the honest one this case insists on: rules the queue owed anyway, tuned two weeks before the industry's naming instead of after; a staff already reading diplomatic-themed mail with narrowed eyes when the campaign class went public; and an adversary whose cost of operating against that ministry, for those nine weeks, was higher — which is all defensive foresight ever gets to claim, and exactly what it is for.

EPILOGUE

Cyber defense compresses this volume's whole argument into its purest economics: the defender cannot prevent, only reprice — every tuned rule, every briefed inbox, every watched flow raises what the intruder must spend. The residue for any SOC: *consume claims at the actor grain, attach dates to owed work, write the revert clause, and grade yourself on the adversary's costs, not on the intrusions that never announce themselves.* And for reading strange inputs anywhere: the coverage objection is real, permanent, and conceded — which is why the only threads worth an hour are the ones whose keepers concede it too, and grade every line anyway.

CASE THIRTY-EIGHT

Planning Paris on the Neutral Pathway: The Rights Desk

A teaching case in sponsorship and media risk — carrying the scenario the industry's planning file excluded

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the framing discipline the record enforces on itself binds every page here: the Manama vote was formally a **partial suspension, never a reinstatement**; the record prices two of the sealed statement's three clauses at **close to zero information** (a sporting power "continuing to do well" is near base rate; the fatigue-and-resolution mechanism is unfalsifiable narrative) and grades neither; and the objections travel up front — the direction may already have been visibly turning (Wimbledon had readmitted neutrals before the seal), the 74–65 margin shows a genuinely divided room a competent planner might have hedged regardless, and a direction-of-travel claim is generous to many outcomes. Every vote, decision, and date is real and cited. The **protagonist and her company are composite characters** — "Céline Marchand" and "Aurora Sports Media" are not real; their rights portfolio and meetings are teaching devices, and no outside-world fact discussed is invented. The case takes **no position** on the eligibility debates themselves — the boycott arguments and the inclusion arguments both had serious advocates, and the classroom inherits the question, not an answer. It takes **no position** on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

WEDNESDAY, SEPTEMBER 6, 2023 — 10:00, PARIS

Céline Marchand ran rights-and-sponsorship risk for Aurora Sports Media — a rights-holder and production partner heading into the busiest cycle of its existence: Paris 2024, eleven months out. Her planning file's hardest question was not logistical but political: **would Russian athletes be at the Games at all?** The industry's working answer, everywhere she looked, was effectively *no*. The International Paralympic Committee's full suspension of NPC Russia had stood since November 2022; flags, anthems, and teams were banned across most Olympic sport; the loud debate of 2023 was whether Paris would exclude

Russians entirely, with boycott discourse live on the other side. Aurora's scenario file mirrored the ambient state: exclusion as base case, everything else a footnote.

The input her intelligence analyst flagged read against the room. A public, timestamped post of **September 4** — from a source with a sealed and publicly graded record, method disclosed as Vedic astrology — stated: *"Russia will continue to do well in international sporting events during this period. More events will start opening up to players from Russia as fatigue sets in from the conflict and sport is increasingly used as a means to help resolve the conflict."*[1]

Marchand's read was a professional's triage of three clauses. *"Continue to do well"* — near base rate for a sporting power; worth nothing. The fatigue-and-resolution *mechanism* — an unfalsifiable narrative; worth nothing. But the operative middle — **"more events will start opening up"** — was a dated direction-of-travel claim about exactly the variable her file had flattened to a footnote, sealed while the prevailing direction was exclusion. And unlike most political questions in her world, this one converted *directly* into production and contract work with long lead times: how neutral-status athletes would be captioned and presented; what commentary guidance would say; how sponsor association clauses handled a Games where Russians competed under conditions nobody had yet defined. Every one of those questions would need an answer before Paris *if* the door moved — and every answer was cheap in September and compressed after any December announcement.

THE THREE POSTURES

Package A — Monitor. Keep exclusion as the base case; treat re-admission as a headline to react to. Cost: zero. Exposure: if the door moved mid-cycle, the entire presentation-policy and contract-language workload would land inside the pre-Games compression, when every calendar in the industry was already full.

Package B — Scenario readiness. Four paper-only workstreams, inside her function's authority, reversible, no capital. *One:* **promote the neutral-pathway scenario** from footnote to carried working case — a memo, a probability, an owner. *Two:* **contract language** — instruct counsel to draft association clauses, neutrality-condition language, and adjustment triggers keyed to athlete-eligibility outcomes into the autumn's renewals, drafted for *some* conditional or neutral status without guessing any specific classification. *Three:* **editorial and presentation policy** — open the production questions now: captioning,

graphics, medal-presentation treatment, commentary guidance, advertising adjacency around such athletes' events. *Four:***sponsor pre-briefing** — walk counterparts through the neutral-pathway branch and the boycott-discourse branch alike, converting any future announcement from fire drill to planned playbook page.

Package C — Repositioning. Bet the slate: commission programming premised on Russian participation, renegotiate deals around it, brief the market. Cost: a portfolio committed to one outcome of a genuinely divided governance process, on one source.

Her general counsel's objection had a fair edge: *"Wimbledon already readmitted neutrals — if the direction is turning, our own trade press will tell us; if it isn't, your astrologer is reading tea leaves about the most divided rooms in sport. Either way, why does this line move budget?"* Marchand's answer, drafted for the planning board: *It doesn't move budget — that's the point. Everything in B is paper, and most of it is work we'd owe in compressed form the week any door opens. What the line changes is the file's* shape*: we currently carry exclusion as if it were settled, and the settled thing is actually a 74-to-65 kind of question* (she would be more right about that number than she knew). *A dated, divergent, falsifiable claim is a good enough reason to make our base case earn its status instead of inheriting it.* The scenario memo needed a probability by Friday — and a named owner.

EXHIBITS

Exhibit 1 — The ambient state, September 2023. IPC full suspension of NPC Russia in force since November 2022; flags/anthems/teams banned across most Olympic sport; 2023's live question: total exclusion from Paris; boycott discourse active; Wimbledon's neutral readmission the one visible counter-signal.

Exhibit 2 — The sealed input. Public post of September 4, 2023, quoted verbatim in the case text. Advisory documentation and grading at jyotishintelligence.com/intel/advisory/IA-RU-039 — the grade rides on the **operative direction clause alone**; the record prices "continue to do well" (near base rate) and the fatigue mechanism (unfalsifiable) at close to zero and grades neither; the same sitting's kept misses named in the file. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The lead-time map (composite teaching device). Contract renewals: drafting through the autumn; retrofit cost after signature: high. Presentation policy: quarters of committee time across production, legal, editorial. Sponsor relationship management: pre-brief cheap, surprise expensive. The pre-Games compression: every calendar full from Q1 2024.

Exhibit 4 — What subsequently happened (public record, for the teaching note). September 29, 2023 — 25 days after the seal: the IPC General Assembly in Manama rejects full suspension of NPC Russia **74–65**, then votes 90–56 for a **partial suspension** — formally still a sanction, never a reinstatement — preserving the neutral pathway to Paris 2024: the first in-window movement of the door, toward opening. **December 8, 2023 — 95 days after the seal:** the IOC Executive Board approves **Individual Neutral Athletes (AINs)** for Paris — Russian passport holders under strict conditions: no flag, no anthem, no team events, no active war supporters. The 2023 planning year's dominant question resolves toward conditional admission; federation-level eligibility cascades through spring 2024 on the IOC's frame.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · Scenario readiness	C · Repositioning
What it adds	nothing	a carried scenario + drafted clauses + opened production questions + briefed sponsors	first-mover programming
What it forecloses	the cheap quarter	analyst and counsel hours	balance across a divided outcome
Cost if exclusion holds	0	paper filed; clauses dormant but harmless	a slate to unwind, publicly
Cost / value if the door opens	the whole workload lands in the compression	December becomes a playbook page	positioned — by luck through a 74–65 room
Speed	—	days to weeks	quarters
Authority required	none	her function's own	board + market
Reversibility	n/a	full	poor
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/broadcast-rights
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — Planning Paris on the Neutral Pathway

Rails: framing exactness is the case's own discipline — **partial suspension, never reinstatement**, in every student telling; two of three sealed clauses are worth ~zero and the record says so first; the Wimbledon counter-signal and the 74–65 margin mean a competent planner *might have hedged unaided* — that deflation is admissible and may win; the eligibility controversy itself is inherited as a question, never adjudicated; no public figures exist for late-planning costs and none may be invented.

LEARNING OBJECTIVES

(1) Direction-of-travel consumption: what a dated directional claim legitimately licenses (re-weighting a scenario file; starting long-lead paper) versus what it cannot (any bet on mechanism, margin, or specific classification) — the clause-triage discipline applied to the loosest claim grade in the volume; (2) the base-case audit: industry planning files inherit ambient states as settled base cases — the divided 74–65 reality behind "settled" exclusion is the case's quiet warning, and a divergent input's chief value is forcing the base case to *earn* its status; (3) lead-time arbitrage in media/rights work: presentation policy, contract clauses, and sponsor relations are quarters-cheap early and compression-expensive late — the September/December asymmetry as the desk's whole economics; (4) barometer reading: a governance vote as a leading indicator of broader normalization pressure — how a rights desk reading the door early also reads the commercial environment early, without ever becoming a policy advocate.

ASSIGNMENT QUESTIONS

1. Write Marchand's scenario memo: the probability she assigns the neutral pathway (defend it against both the Wimbledon signal and the exclusion consensus), the owner, the four workstreams sequenced, and the general counsel's question answered in one paragraph.
2. Run the clause triage yourself: three clauses, two priced at zero by the record's own file. Is the record's pricing right? Construct the strongest case that even the operative clause is "generous to many outcomes" — then state what, exactly, a claim must add beyond direction to earn more than paper-only responses.
3. The 74–65 vote: what does a nine-vote margin in a 139-member room teach about the word "settled" in any industry's planning file? Design the standing audit — how often, by whom — that keeps a scenario file's base case honest about the votes

behind it.

4. Draft the presentation-policy question list in full (captioning, graphics, medals, commentary, adjacency) — and mark which items could be answered generically in September versus which genuinely required December's specific conditions. What fraction of the workload was date-independent? (That fraction is the case's value, measured.)
5. Complete Exhibit 5 and sign it; fill the hindsight row as of December 9, 2023, and honor the deflations — the room was divided, the signal existed, and the desk might have hedged anyway.
6. **The machine staff-officer.** An AI governance-watch system now tracks every federation's votes, statements, and member alignments, and produces continuous door-direction probabilities for every excluded delegation across every sport. The scenario file re-weights itself. What remains human — the presentation-policy judgments? the sponsor conversations? the decision that a base case has stopped earning its status? Write the desk's rule for consuming machine door-probabilities, including who owns the memo when the machine and the ambient consensus disagree.

THE BOARD PLAN (75 MIN)

Open on Exhibit 1 and the planning file (10): exclusion as inherited base case — have students state what "settled" means and what evidence they have for it (little). The input and clause triage (15): Exhibit 2 with Q2 run immediately — two clauses die on the record's own pricing, and the room should watch the record do the killing. The lead-time map (15): Exhibit 3, and Q4's date-independent fraction — the case's economics made concrete. Commit (15): memo + Exhibit 5. Reveal Exhibit 4 (10): Manama at 25 days — 74–65, partial suspension, framing kept exact — then AIN at 95; run Q3 on the margin. Close (10) with the barometer point and its boundary: reading a door early is analysis; cheering a door is advocacy — the desk does the first, only.

WHAT HAPPENED, AND THE GRADE

The door moved on the sealed direction, twice, inside the window. Twenty-five days after the seal, the IPC General Assembly rejected full suspension 74–65 and imposed a partial suspension — a sanction still, and this case keeps that framing exact — preserving the neutral pathway; ninety-five days after the seal, the IOC approved Individual Neutral Athletes for Paris under strict conditions, resolving the planning year's dominant question toward conditional admission and setting the frame the federations cascaded through spring. The record grades the call on its **operative direction clause alone**, prices the other two clauses at approximately zero and grades neither, and prints the deflations this note enforces — the visible Wimbledon turn, the divided room, the generosity of directional language. For the composite desk, Package B's ledger was the lead-time arbitrage realized: a scenario carried from September, clauses drafted before signature season closed, production questions opened with a

quarter of runway, sponsors for whom December 8 was a briefing they had already had — against an industry that met the same announcement inside the compression, re-planning at each federation's turn.

EPILOGUE

This case brings the volume to the desk least like a war room and proves the disciplines travel intact: clause triage, base-case audits, lead-time arbitrage, the scenario memo with an owner. Sport, media, and sponsorship live downstream of governance votes the way ports live downstream of weather — and the residue is the same on both frontages: *the ambient state is not a base case; it is a candidate. Make it show its margin — and start the paper the day a dated, divergent input gives you a reason.*

CASE THIRTY-NINE

Retasking Before the Anniversary: The Collection Manager

A teaching case in intelligence collection — pointing assets that already exist, and why a clean negative is a product

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and the load-bearing constraint is stated first and enforced throughout: the kinetic window in this case is graded **PARTIAL and remains UNRESOLVED on public evidence.** This case does **not** argue that any nuclear test occurred, did not occur, or that any mechanism or supplier attribution was confirmed — as of the window's close, no public reporting (IAEA, commercial-satellite analysts, arms-control researchers) had tied any named coordinate to an observed test signature, and that absence is part of the record. The record's own concessions travel too: zero independent public chatter tied the coordinates to a test; and a March 2025 antecedent had already put a comparable supplier rumor into circulation, so the later Russian statement may be a reaction to ambient rumor rather than to this seal. The supplier-attribution narrative that Russian state organs later layered onto the window is **not endorsed, repeated as fact, or evaluated** here — it is treated only as a dated information-operations event. Every statement, vote, and disclosure is real and cited. The **protagonist is a composite character** — "Elena Sørensen" is not a real person; no actual collection agency's operations are depicted, and no outside-world fact discussed is invented. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

MONDAY, JANUARY 5, 2026 — 08:00, A NATIONAL COLLECTION-MANAGEMENT DESK

Elena Sørensen managed a tasking queue — the desk that decided which of a finite set of *already-existing* overhead-imagery passes and seismic/radionuclide review priorities got pointed where this week. Her authority was specific and bounded: she could **re-prioritize existing looks**, not commission new satellites, open new stations, or set policy. That boundary is the whole reason this case has a cheap answer.

The item her open-source cell had flagged came in two forms, days apart. A **private** advisory shared with agencies on January 1, and an identical **public** bulletin on January 4 — from a source with a sealed, publicly graded record, method disclosed as Vedic astrology — naming a probable Ukraine-zone nuclear-test event in a **March 10–15, 2026 window**, at a named site: the Chernobyl Exclusion Zone.[1] Within seventy-two hours of the seal, two unaffiliated public figures had independently echoed a Ukraine-nuclear-test-with-Western-help scenario on a similar timeline — noise, but dated noise.

Sørensen's read was a collection manager's, and it turned on one fact about her own assets: **the Chernobyl Exclusion Zone was not virgin territory.** It sat inside an elevated radiological and safeguards-monitoring baseline from the 1986 accident and the post-2022 occupation disputes; it was already imaged on a recurring cycle by national and commercial systems; and it already fell inside the international seismic and radionuclide network built for test-ban monitoring. The named site was, in other words, *already being looked at.* The decision the bulletin posed was not whether to build collection — it was whether to move three named boxes to the **top** of a queue whose assets were already pointed at roughly the right region.

She was rigorous about what she would and would not credit. The supplier-attribution story that might attach to such a window — which state helped whom with what warhead — was policy narrative she had no basis to evaluate and no need to. What she could act on was narrower and cleaner: a dated, public, falsifiable **site-and-window** claim, against a region her assets already covered. And the product she could produce was symmetric in a way most of her taskings were not: whether the window opened with a detectable signature **or passed clean**, an independent look would yield a data point — and a documented *negative* against named coordinates might be worth as much as a detection, because it could constrain a narrative that a bare claim could not.

THE THREE POSTURES

Package A — Monitor. Leave the standing revisit cadence as-is; let automated thresholds catch anything. Cost: zero. Exposure: if the window drew attention — from any direction — her chain of command would hold no *independent, deliberately-tasked* read, only the ambient claim and whatever an automated alert happened to trip.

Package B — Priority retasking. Three actions, inside her existing authority, executable in the standing cycle. *One:* move the named Exclusion-Zone boxes to the top of the existing overhead-imagery tasking queue — a same-day-to-72-hour reallocation under standard revisit cadences, not a new acquisition. *Two:* flag the relevant seismic and radionuclide stations for **manual waveform review** ahead of automated alert thresholds — a review-priority change, essentially free. *Three:* set the collection to run **through the window**, so that a clean negative is captured and documented as deliberately as a positive would be. All of it bounded to the site-and-window claim; none of it requiring, or endorsing, any attribution narrative.

Package C — Elevated posture. Escalate to a new-collection request, cross-agency alerting, and a policy-level warning built on the window. Cost: procurement clocks and sign-offs that dwarf the lead time, spent against a call graded PARTIAL with the kinetic vector unresolved — and the risk of a collection agency appearing to endorse an attribution narrative it cannot verify.

Her section chief's objection was the arms-control desk's reflex: *"Russian nuclear rhetoric is constant, and this one comes with an astrologer attached. If we retask against it and it's noise, we've spent looks — and if word gets out we tasked on a horoscope, we've spent credibility."* Sørensen's answer, logged in the tasking justification: *We're not tasking on the attribution story — I can't evaluate that and I'm not trying to. We're tasking a named site inside our existing coverage, for a bounded window, because the marginal cost of moving three boxes up a queue that's already pointed there is close to zero, and because a documented look — positive or negative — is a product my chain of command doesn't currently have. If the window passes clean, that clean negative is the most useful thing this desk could hand a policy office arguing against an escalation. The looks already exist. The only question is priority.* The window opened in sixty-five days; the retasking decision was hers this week.

EXHIBITS

Exhibit 1 — The existing coverage baseline. The Chernobyl Exclusion Zone under elevated radiological/safeguards monitoring since 1986; recurring national and commercial imagery cadence; inside the international seismic/radionuclide network. The collection manager's authority: re-prioritize existing looks (days); not new acquisition (months–years) or policy.

Exhibit 2 — The sealed input. Private advisory January 1, 2026; identical public bulletin January 4 (multi-language): probable Ukraine-zone nuclear-test event, March 10–15, 2026 window, Chernobyl Exclusion Zone site. Advisory documentation at jyotishintelligence.com/intel/advisory/IA-RU-020 — **kinetic window graded PARTIAL / UNRESOLVED**; the record's own verification note records **zero independent public chatter** tying the coordinates to a test, and flags the March 2025 antecedent that may explain the later Russian statement independently. Ledger and verification as in prior cases (snapshot of 2 Aug 2026: 106 graded, 82/10/5/9, Brier 0.0903, misses permanent).

Exhibit 3 — The retasking economics (composite teaching device). Overhead reallocation: same-day–72h under standard revisit. Manual waveform-review flag: essentially free. New satellite / new station / new authority: months–years — outside the desk's power and irrelevant to the decision. Lead available: 65 days (public seal to window); 10 days (later coordinate disclosure to window). The lead was never the constraint.

Exhibit 4 — What subsequently happened (public record, for the teaching note — dated events, NOT endorsements of any claim's content). Within 72h of the seal: two unaffiliated public figures echo a Ukraine-nuclear-test scenario. **February 24, 2026 (fourth invasion anniversary):** Russia's SVR issues an official statement alleging Western-assisted nuclear transfer, naming a specific warhead and delivery vehicle — a claim Ukraine and a UK official denied outright. Within hours/days: escalation via a senior security-council figure; a State Duma foreign-affairs appeal (Feb 25); full Duma adoption (Feb 26) addressed to four Western legislatures/bodies. February 28: three specific Exclusion-Zone coordinates (plus one Iran coordinate) disclosed through dated government channels including US Army Counterintelligence's iSalute intake. **Through the window's close: no public reporting from any independent body tied any coordinate to an observed test signature** — the kinetic window stands UNRESOLVED. What is *documented* is a four-organ Russian information-operations cascade built on a site-and-window that had been on the public record 51 days before the SVR statement.

Exhibit 5 — The Decision Package, priced (recommendation left to the student).

	A · Monitor	B · Priority retasking	C · Elevated posture

What it adds	nothing beyond ambient alerts	a deliberate, documented look — positive or negative — over the window	new-collection request; policy warning
What it forecloses	an independent read	three boxes' worth of queue priority	credibility if it reads as endorsing attribution
Cost if the window is noise	0	a few reprioritized looks; a clean negative on file	procurement clocks + a spent narrative
Cost / value if activity appears	only what automation happened to catch	an independently-tasked detection, in hand	over-committed on a PARTIAL call
Value of a clean negative	none captured deliberately	a product to constrain the narrative in-cycle	disproportionate to the ask
Speed	—	standing revisit cycle (days)	months
Authority required	none	collection manager's own	policy office + acquisition
Reversibility	n/a	full	poor
Recommendation & rationale		*(yours)*	
In hindsight	*(the window is UNRESOLVED — write yours anyway)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/collection-window
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — Retasking Before the Anniversary

Rails — absolute: the kinetic window is **UNRESOLVED and stays so** — no telling asserts a test happened or didn't; the supplier-attribution narrative is **never endorsed, repeated as fact, or adjudicated** — it is a dated information-operations event, full stop; the record's own concessions (zero corroborating chatter; the March 2025 antecedent) lead; the boundary between *tasking a site* and *crediting an attribution* is the case's spine and must never blur; the case is not about nuclear weapons policy — it is about collection-queue economics.

LEARNING OBJECTIVES

(1) The existing-asset distinction: the entire cost structure of a collection decision depends on whether the asset must be *built* or merely *re-pointed* — retasking coverage that already exists is the cheapest intelligence action there is, and the value math is nothing like a new-architecture decision; (2) the symmetric product: a deliberately-tasked **clean negative** is a product, often as valuable as a detection — especially against an active narrative — and desks that only value positive findings systematically under-task; (3) tasking-the-claim versus crediting-the-narrative: how to act on a dated site-and-window input while refusing an attribution story you cannot evaluate — the sharpest separation-of-concerns exercise in the volume; (4) grading the unresolved: what it means to make and defend a decision on a call that never resolves — the PARTIAL/UNRESOLVED status as a permanent condition, and why the decision's quality is independent of an outcome that never arrives.

ASSIGNMENT QUESTIONS

1. Write Sørensen's tasking justification: the three actions, the explicit statement that the desk is tasking a site and not crediting an attribution, and the sentence defending the value of a clean negative to a policy office. It must survive a leak that says "agency tasked collection on an astrologer's forecast."
2. The separation-of-concerns problem: draw the exact line between "task the named coordinates" (defensible) and "brief that a Western-assisted test is likely" (not). Where on that line does a collection manager stop, and what must never appear in the tasking record?
3. The clean-negative argument, developed: why is a documented negative against named coordinates worth more *during an active information-operations cascade* than in a quiet period? Price it against the do-nothing baseline, where the only counter

to a narrative is a bare denial.

4. This call never resolved. Construct the framework for evaluating a *decision* whose *outcome* is permanently unknown — what makes Package B right or wrong independent of whether a test ever occurred? Connect to the prevention-class cases (15, 33): three flavors of the unscoreable, three different disciplines.
5. Complete Exhibit 5 and sign it; the hindsight row cannot be filled from an outcome — write instead what the desk *would hold* at the window's close under each package, and which holding best serves a chain of command facing a live diplomatic cascade.
6. **The machine staff-officer.** An AI collection-management system continuously scores every named site against every dated claim and auto-proposes retasking — including against claims from sources analysts distrust, and including the low-glamour clean-negative tasks humans under-prioritize. What does it get right (the existing-asset math; the symmetric value of negatives) and where must a human hold the line (the tasking-versus-crediting boundary; what enters the record; what gets briefed as "likely")? Write the collection desk's governance rule for AI-proposed tasking on unverified single-source claims.

THE BOARD PLAN (75–90 MIN)

Open on Exhibit 1's baseline (15): the class must internalize that the site was *already covered* — the whole case collapses to "queue priority," and students who miss this will over-price the decision. The input, bounded hard (15): Exhibit 2 with the UNRESOLVED grade and both concessions read first; then Q2's separation line drawn on the board — tasking versus crediting — as the session's discipline. The clean-negative insight (15): Q3 — most students value only detections; the case exists partly to fix that. Commit (15): tasking justification + Exhibit 5. Reveal Exhibit 4 (15): the anniversary cascade, dated and labeled strictly as information operations, the coordinates, and the window closing UNRESOLVED — with the rail restated every few sentences: nothing here says a test happened. Close (15) on Q4: judging a decision whose outcome never comes, and the three-way comparison with the volume's other unscoreables.

WHAT HAPPENED, AND THE STANDING

The window opened and closed without any independent body — the IAEA, commercial-satellite analysts, the arms-control research community — tighing any named coordinate to an observed test signature; the record holds the kinetic window **PARTIAL and UNRESOLVED**, and this case asserts nothing past that. What *is* documented, and dated, is a different thing: a four-organ Russian information-operations cascade — an SVR statement, a senior official's threats, a two-day State Duma appeal to four Western legislatures — built on a site-and-window that had sat on the public record fifty-one days before the SVR spoke, with the record's own file noting zero corroborating chatter and a prior-year antecedent that may explain the cascade independently. The teaching value lives entirely in the decision, not the

outcome: a collection manager who retasked existing looks against named coordinates for a bounded window spent almost nothing and would have held, at the window's close, the one product a policy office facing that cascade did not otherwise have — a deliberate, documented read, and most likely a clean negative. A desk that judged the decision by waiting for the window to "resolve" is still waiting, which is precisely the condition the discipline must survive.

EPILOGUE

This case closes the volume's intelligence-and-security arc on its subtlest note: the cheapest action in the book — moving three boxes up a queue already pointed at the right place — attached to the least resolvable outcome in the book. Together they isolate the thing this volume has argued from case one: **decision quality is not outcome quality.** A clean negative nobody remembers needing, produced deliberately, during a narrative that a bare denial could not touch, is foresight doing its quietest and most honest work — and a desk that can task for it, refuse the attribution it cannot verify, and defend both in a leak-proof record has learned everything the strange inputs in this book were ever able to teach.

CASE FORTY

The Black Sea Test: The Standing Intake Lane

A capstone case in intake design — whether a theatre's whole year of sealed calls justifies a lane, tested on the record's own base rate

AUTHORSHIP DISCLOSURE. This case was written by **Vijay Jyotish**, operator of the public forecast record referenced in Exhibit 2, published at jyotishintelligence.com; the author has a direct interest in that record — and this capstone concedes its own strongest weakness before anything else: the cluster is **one theatre, one year, four of five calls from a single sealed sitting** — not a random sample, and it does not claim to be. Two of the five carry conceded gaps between maximal claim and certified grade (the storm was a scenario-class match, not the sabotage the text framed; the bridge campaign's maximal "beyond repair" end-state never occurred), and one is graded **PARTIAL** with two of four named vectors dead — all carried at full weight, none smoothed into the aggregate. The cluster's 4-HIT/1-PARTIAL record is **consistent with, not dramatically better than**, the theatre-wide base rate, and this case argues that consistency is the point, not a lucky year. Every seal, grade, and date is real and cited; four of these five are the underlying files of cases 7, 10, and others in this very volume. The **protagonist is a composite character** — "Commodore Anders Holt" is not a real person; no actual command's intake design is depicted, and no outside-world fact discussed is invented. The case takes **no position** on whether the forecast method referenced in Exhibit 2 works.

METHOD DISCLOSURE. The record referenced in Exhibit 2 discloses its input method plainly: **Vedic astrology.** The case treats the record as a dated document and leaves the epistemic question to the classroom.

WEDNESDAY, JANUARY 3, 2024 — 09:00, AN ALLIED MARITIME COMMAND

Commodore Anders Holt did not own a desk in the sense the other thirty-nine cases in this volume mean it. He owned the **design of intake** for a maritime theatre's warning cycle — the standing orders that determined how a claim, any claim, got triaged before it ever reached the officers who would act on it. His question was never "is this input right?" It was the question one level up: *what standing rule decides which inputs my people even see?*

The occasion for the question was an after-action his J2 cell had assembled, covering the previous year in the Black Sea. It documented something that stopped him. A single unaccredited public source — method disclosed, in its

own materials, as Vedic astrology — had sealed, between April and September 2023 and mostly in one 59-post sitting, **five dated, falsifiable claims** about his theatre. Over the following year all five had materialised across five entirely different desks, none of which sat in the same meeting in the ordinary course of business:

- an **ISR collection manager** (a doctrine-level read of raiding intent, sealed before any indicator) — HIT; - a **port-resilience planner** (a destructive sea event against the Crimean coast and fleet) — HIT, scenario-class, mechanism caveat conceded; - a **sanctions-enforcement officer** (an all-out campaign against the Kerch bridge) — HIT, maximal end-state caveat conceded; - a **fleet-dispersal officer** (elevated holiday-season alert against a lull) — HIT; - a **war-risk underwriter** (high-seas sabotage vectors against Russia-serving vessels) — PARTIAL, two of four named vectors dead.

Four HIT, one PARTIAL. And — the fact Holt kept returning to — **not one of those five desks would ever have compared notes with another.** An ISR tasker and a marine underwriter inhabit different institutions, let alone different meetings. Each call, in the year as it actually ran, either reached the right desk by some desk's own ad hoc judgment about an unaccredited source, or reached no one.

The question the after-action forced was institutional, and it was Holt's alone to answer: *does a full theatre-year with this range and hit rate justify a standing, low-friction lane — a rule that logs and routes dated, sealed, checkable claims on their own merits rather than on their source's pedigree — or is one lucky sitting, mined for pattern after the fact, exactly the kind of thing a serious command should refuse to build doctrine on?*

THE THREE POSTURES

Package A — No lane. Keep the status quo: unaccredited claims filtered by whatever ad hoc pedigree judgment each receiving officer applies, if one reaches them at all. Cost: zero. This is precisely the year the after-action described — five claims, mostly one sitting, reaching the right desk by chance or not at all.

Package B — The standing lane. One paragraph in a J2 standing order: a dated, sealed, falsifiable claim meeting a minimal specificity bar gets **logged, timestamped on receipt, and routed** to whichever desk it names — with every routing and every subsequent action decision remaining entirely with the receiving desk. The lane asks no one to believe anything; it converts "twenty

minutes of triage, or never" into "twenty minutes of triage, before." Cost for the whole year: duty-officer staff-minutes per claim, borne below the decision-maker level. Fully reversible if it stops earning its keep.

Package C — The credence engine. Go further: stand up a cell to score, weight, and brief unaccredited claims as graded intelligence products feeding the theatre estimate. Cost: real headcount, a standing commitment to adjudicate sources the command cannot audit, and the institutional hazard of a warning cell that has taken ownership of others' forecasts.

Holt's own strongest objection was the one his intelligence director raised: *"One theatre, one year, four calls from a single sitting — that's not a sample, it's an anecdote. Build doctrine on it and you've enshrined survivorship."* Holt's working answer, the one this case exists to test: *The objection would be fatal if I were claiming this year was unusually good. I'm not — and the record won't let me. The theatre-wide ledger this cluster is drawn from runs 106 graded calls near a 0.09 Brier with every miss kept, and the cluster's four-and-one is consistent with that base rate, not better than it. That's the case for a lane, not against one: a claim class that performs near its own long-run average, routed for staff-minutes, is exactly what a cheap standing rule is built to catch. I'm not pricing a lucky year. I'm pricing a stream.* Whether that answer held — whether a base-rate-consistent cluster is a fair test of a general rule, or a rationalization of a small sample — was the decision the standing order turned on.

EXHIBITS

Exhibit 1 — The five calls (the record's own grades).

Asset	Sealed	Lead	Desk it named	Grade
IA-RU-004	30 Apr 2023	116 d	ISR collection manager	HIT
IA-RU-006	4 Sep 2023	83 d	port-resilience planner	HIT (scenario-class; mechanism caveat)
IA-RU-019	4 Sep 2023	268 d	sanctions-enforcement officer	HIT (maximal end-state caveat)
IA-RU-023	4 Sep 2023	113 d	fleet-dispersal officer	HIT
IA-RU-030	4 Sep 2023	353 d	war-risk underwriter	PARTIAL (2 of 4 vectors dead)

Four from a single 4 September 2023 sitting; IA-RU-004 sealed earlier. Each graded independently against a reader and decision fixed **before** any combined framing existed.

Exhibit 2 — The record behind the cluster. Theatre-wide ledger as of 2 August 2026: 106 graded (82 HIT, 10 NEAR, 5 PARTIAL, 9 MISS), mean Brier 0.0903, misses retained at full weight, every seal OpenTimestamps-anchored before its event, every statistic recomputable from the published dataset. Cluster and per-asset documentation at jyotishintelligence.com/intel/advisory/ (IA-RU-004, -006, -019, -023, -030). The cluster's 4/1 is **consistent with** this base rate — the case's spine.

Exhibit 3 — The lane, priced (composite teaching device). Trigger: a dated, sealed, falsifiable claim meeting a minimal specificity bar. Action: log, timestamp-on-receipt, route to the named desk. Cost: duty-officer minutes per claim. Non-cost: no belief, no adjudication, no action obligation — every decision stays with the receiving desk. Reversibility: delete the paragraph.

Exhibit 4 — The Decision Package, priced (recommendation left to the student).

	A · No lane	B · Standing lane	C · Credence engine
What it adds	nothing	dated claims logged and routed on merit	scored products in the estimate
What it forecloses	routing that doesn't depend on chance	duty-officer minutes	the command owning others' forecasts
Cost if the stream is mostly noise	0	staff-minutes triaging noise	a cell adjudicating the unauditable
Cost / value if the stream carries signal	signal reaches desks by luck or not at all	signal routed cheaply, every time	signal over-processed, ownership assumed
Speed	—	days (one standing-order paragraph)	months (headcount + doctrine)
Authority required	none	J2 standing-order authority	command + resourcing
Reversibility	n/a	full (delete the paragraph)	poor
Recommendation & rationale		*(yours)*	
In hindsight	*(filled at grading)*		

PRINTABLE CASE HANDOUT — FOR YOUR STUDENTS
jyotishintelligence.com/cases/intake-lane
hosted on the record's site and deposited at Zenodo · given away on purpose

TEACHING NOTE · INSTRUCTORS ONLY — THE CASE BODY ENDS AT THE DECISION POINT; THE ANSWER LIVES HERE

Teaching Note — The Black Sea Test

Rails: the small-sample objection is the case's own strongest voice and may win — one theatre, one year, mostly one sitting, and the note must never let the aggregate bury the caveats; the two conceded gaps (RU-006 scenario-not-mechanism; RU-019 maximal-end-state-unmet) and the RU-030 PARTIAL stay at full weight; the argument is **base-rate consistency, not an exceptional year** — any telling that sells the cluster as unusually good has inverted it; this is an intake-*design* case, one level above any single desk's action, and the classroom decides a standing rule, not a call.

LEARNING OBJECTIVES

(1) The pedigree gate versus the performance gate: intake systems built to filter *individual inputs by source* cannot price a *stream by record* — and the whole volume's recurring "which desk, at what grain" question resolves, at the design level, into whether a command admits claims on merit or on provenance; (2) range versus accuracy: five hits across five unconnected desks is a claim about *reach*, not about being right on any one decision — students learn to state precisely what a cross-desk record does and does not establish; (3) base-rate-consistency as an argument *for* a cheap rule: a claim class performing near its own long-run average, routed for staff-minutes, is the exact profile a low-cost standing lane is designed to capture — the counterintuitive core; (4) survivorship self-defense: how to test one's own cluster for curation (independent prior grading, timestamped sequence, the theatre-wide denominator) and concede honestly where the test is weak.

ASSIGNMENT QUESTIONS

1. Write Holt's J2 standing-order paragraph: the specificity/dating bar, the log-timestamp-route mechanism, the explicit non-obligation of the receiving desk, and the reversibility clause. One paragraph — the constraint is the point.
2. The survivorship objection, adjudicated: is a base-rate-consistent cluster a *fair test* of a general intake rule, or a rationalization of a small sample? Build both cases fully, using Exhibit 2's theatre-wide denominator, and reach a defensible verdict.
3. Pedigree gate versus performance gate: for each of the five desks, state what pedigree check it would ordinarily apply to an unaccredited source and why that check fails to price a stream. Then design the minimal specificity bar a performance gate would use instead — and stress-test it against a flood of low-quality claims.

4. Range is not accuracy. Write the two-paragraph honest statement of exactly what "four hits across five unconnected desks in one theatre-year" establishes and what it does not — the statement Holt would put above his standing order so no one over-reads it.
5. Complete Exhibit 4 and sign it. Then reflect: this is the volume's fortieth case, and it turns the book's method on itself — after thirty-nine desks, does the standing lane clear? Write the recommendation you would defend to a command that has read all forty.
6. **The machine staff-officer.** An AI intake system could run the lane at zero marginal cost — logging, timestamping, routing every dated public claim on earth to whichever desk it names, and scoring each source's live record automatically. That collapses Package B's cost to zero and makes Package C's adjudication cheap. What does that change about Holt's decision, what new failure mode does it create (a routing firehose; automated credence the command implicitly owns), and where must the human standing-order authority still draw the line? Write the design rule — the last machine-staff-officer question in the book, and the one that governs all the others.

THE BOARD PLAN (90 MIN — THE VOLUME'S CAPSTONE; TEACH IT LAST)

Open by naming the shift (10): thirty-nine cases asked "what should this desk do?"; this one asks "what should the command let desks see?" — the level change is the lesson. Exhibit 1 and the five-desk range (15): the fact that no two readers would ever meet is the case's engine — draw it out. The survivorship objection at full strength (20): Q2 — let the room try to kill the cluster, then introduce Exhibit 2's base-rate denominator and watch the objection transform into the argument. Pedigree versus performance gate (15): Q3, the transferable design distinction. The standing order (15): Q1 written, one paragraph, plus Q4's over-read guard. Commit (10): Exhibit 4 signed. Close (5) on Q5: forty cases in, does the lane clear? — and let the book end on the students' verdict, not the author's.

WHAT HAPPENED, AND THE STANDING

There is no single outcome to reveal here, because the case's subject is a rule, not a call — and that is the capstone's final discipline. What the record shows is that all five calls materialised across their five desks at the grades in Exhibit 1: four HIT, one PARTIAL, two of the hits carrying conceded gaps between maximal claim and certified grade, all drawn from a theatre-wide ledger of 106 graded calls near a 0.09 Brier with every miss kept and every seal anchored before its event. The cluster is consistent with that base rate, not better than it — and the case for the lane rests entirely on that consistency: a claim class that performs near its long-run average, routable for staff-minutes, either reaches five unconnected desks through a one-paragraph standing rule or reaches them by the chance this after-action documented. The record concedes it is one theatre, one year, mostly one sitting — not a random

sample — and lets the reader weigh whether that is a fair test of a general rule. This note, and this book, decline to answer for the reader. The whole volume has been the argument; the standing order is the student's to sign.

EPILOGUE

This is the fortieth case and the only one that turns the method on the institution that would use it. Every prior case handed a desk a strange input and asked what it made cheap; this one hands a command forty desks' worth of evidence and asks whether the cheapest possible rule — *log it, timestamp it, route it, oblige no one* — is worth a paragraph. The book's answer, offered once and then left to the room: a pedigree gate protects a command from bad sources and, in the same motion, from the dated, checkable, occasionally decisive claim that arrives wearing the wrong credentials. A performance gate, priced in staff-minutes and reversible by deletion, risks only the minutes. Forty desks in, that is the whole of what this record can teach an institution — and the last thing it asks is not belief, but a standing order, and the discipline to grade what comes through it exactly as ruthlessly as this record grades itself.

TEMPLATE 1 OF 6 · THE TEMPLATES

The Decision Package

> **These printed forms are representative.** They are sized to show each template's structure on a book page — not to be written on. The **fully functional write-in versions**, one template per page with real working space, are free and permanent at **jyotishintelligence.com/cases/templates** and deposited at Zenodo. Print those to use; these pages are here so you know what each instrument looks like and when to reach for it.

The commitment instrument every board plan references. One page; photocopy or extract. A package without an owner and an expiry is a memo.

THE INPUT (verbatim, with source and date)	
Anteriority checked?	yes / no — if no, stop
Decision owner	
Decision expiry (date or trigger)	

	A · WATCH	B · QUIET READINESS	C · DECLARED POSTURE
What it adds			
What it forecloses			
Cost if nothing happens			
Cost / value if it happens			
Speed to execute			
Authority required			
Reversibility			

Recommendation & rationale			
In hindsight (filled at grading)			

Audit rule: read the C column first — if it is thin, the menu is rigged. · Grading rule: the hindsight row is filled on the expiry date, whether or not anything happened. · Classroom rule: every student signs before the room learns the outcome.

ALL TEMPLATES, FULLY FUNCTIONAL & PRINTABLE —
jyotishintelligence.com/cases/templates
hosted on the record's site and deposited at Zenodo

TEMPLATE 2 OF 6

The Documented Hold

The four-line artifact that separates a judged non-action from silent inaction. Used in the Xi Weekend, Tariff Shock, and City-desk cases; usable at any desk, any week.

THE INPUT (verbatim, source, date)	
What we assessed (one sentence)	
What we judged, and the dated reason	
REVISIT TRIGGER (observable; a duty officer could call it)	
NAMED FALSIFIER (what would prove this hold wrong)	
Owner / date / expiry	

Discipline check: if the revisit trigger requires judgment to recognize, rewrite it until it doesn't. · A hold with no falsifier is a belief, not a decision. · File it where the five-year lookback will find it.

TEMPLATE 3 OF 6

The Escalation Ladder

Rungs from documented inaction to full activation, priced in both worlds before either is known. Used in the Xi Weekend and Moscow cases.

Rung	Action	Cost now	Cost / value if the event is real	Rev.?	Authority
0	do nothing, log nothing				
1	documented hold (see template)				
2	convene / brief				
3	operational hedges				
4	financial hedges				
5	escalate to board / principals				

Worst worst-case rung: _____ · **Worst expected-case rung:** _____

Rules: price the credibility currency, not just money — every convening spends alarm credibility. · Mark the first rung that cannot be climbed down from. · The cheap middle is usually rung 1–2; if your analysis lands on 0 or 5, check the menu for rigging.

TEMPLATE 4 OF 6

The Tripwire Sheet

Pre-committed conditional movement: written thresholds that convert watchfulness into action without a 2 a.m. committee. Used in the Moscow and Storm cases.

#	Trigger (observable, unambiguous)	Threshold	Owner	Pre-authorized action	Rung
1					
2					
3					
4					
5					

Pre-authorization signed by: _____ **Date:** _____ **Review date:** _____

Discipline check: could the night-duty officer execute each row alone, from its own words? If not, judgment has re-entered — rewrite. · Rehearse once before the window: find the seized bollard while finding it is cheap.

TEMPLATE 5 OF 6

The Intake Note

The one-paragraph discipline for an input that can be neither believed nor ignored. Neither over-reads nor under-reads; used in every case in this volume.

THE INPUT (verbatim, source, date, seal/timestamp link)	
What it asserts (only what the words support)	
What it does NOT assert (the inference you are tempted to add — name it as yours)	
Class and level (event / window / axis / class-only)	
Falsifiable? How, and by when?	
Cheapest hedge it points at, and its price	
Disposition (log · hedge · escalate) + **revisit trigger**	

Rule: the seam between what a warning says and what the reader infers is where analysis becomes contamination — own the inference explicitly or strike it. · Twenty minutes, not zero; log it either way.

TEMPLATE 6 OF 6

The Verification Walk

The student exercise this book is built to survive: check the case's central artifact yourself, tonight, from a phone. Assign before or after the session — before is braver.

Step	What to do	What I found
1 · Anteriority	Open the sealed artifact link in the case's exhibits. Read the platform timestamp. Did the text exist before the event?	
2 · Integrity	Run the verifier (jyotishintelligence.com/verify). Does the hash match? Could the text have been altered after the fact?	
3 · Grading	Find the advisory's grade on the ledger. What did the record concede against itself (misses, partials, coincidences, refusals)?	
4 · Outcomes	Check two load-bearing case facts against sources independent of the record's operator. Do they hold?	
5 · Verdict	What survives all four steps — and what would you now weight differently in the case discussion?	

Rule: the author's interest is disclosed on page one of every case — which is exactly why nothing in this walk requires trusting him.

APPENDIX

The Machine in the Loop

How to run this book's discipline with AI decision support — and what every teaching note's final question has been building toward.

Every case in this volume ends its assignment questions the same way: a machine staff-officer arrives, and the class must decide what it may fill, what it must never sign, and how it is audited. That repetition is deliberate. Decision support is the fastest-moving fact about every desk these cases depict, and a casebook that ignored it would be teaching the last decade. The pattern that emerges across the twenty-four questions is stable enough to state once, here, as doctrine.

What the machine does better, starting now. Intake triage at scale — the twenty-minute rule run on everything, so no dated, falsifiable input dies unread in a feed. Drafting: Decision Packages, escalation ladders, continuity menus, screening packages, generated in minutes with the pricing rows pre-filled for human attack. Watch: tripwire states, round cadence, trade profiles, founder sentiment, computed continuously from sources no human cell can cover. Memory: every hold, every falsifier, every expiry, resurfaced on its date instead of when someone remembers.

What the machine inherits rather than fixes. The room's reflexes. A model trained on consensus commentary prices the third shape low, rates the narrow floor unlikely, and recommends de-risking at the panic's peak — with beautiful prose (the Tariff Shock, the Fifteenth Ballot, the Third Shape). An accuracy-rewarded scorer converges on consensus-followers (the City Hall Window). A pattern-finder rediscovers scheduling coincidences as signal (Two Readers). Machines do not launder the epistemics of this book's inputs; they industrialize whatever epistemics they are given.

What stays human, in every case, without exception. The signature. The Decision Package's recommendation row and its hindsight row; the hold's falsifier; the tripwire's pre-authorization; the designation; the go. Not because machines cannot emit those strings — because accountability is what those cells *are*, and an unsigned package is a memo whoever drafted it. The record

behind this book runs the same rule on itself: the machine can hash, anchor, recompute, and display, and a named human owns every grade and every miss.

The audit, stated once. Whatever the machine fills gets the same three questions this book asks of every strange input: is it dated, is it falsifiable, and what did it concede against itself? A decision-support system that cannot show its priors, keep its misses, and survive a regrade is a consensus engine wearing instruments — and the whole apparatus of this volume, from the seals to the ledger to the blank form in the previous appendix, is the working example of what passing that audit looks like.

Teach the cases with the machine in the room. Let it draft; make the students attack the C column it wrote. Let it argue the opposite posture; make them price its confidence against its training data. And end where every teaching note ends: the machine fills the cells it can defend, the human signs, and reality — as always in this book — referees.

APPARATUS

The Cases, Their Files, and How to Check Everything

The map from each case to its underlying sealed record, the grades as the record itself assigned them, and the verification path a student can walk tonight.

THE CASE-TO-FILE MAP

Every case in this volume stands on a sealed public artifact. The table below is the spine of the book: what was sealed, when, what arrived, and what grade the record itself assigned — including the grades that went against it. Dates are the record's own; each advisory page carries the full sealed text, the provenance chain, and the grading rationale, at jyotishintelligence.com/intel/advisory/(file).

#	Case	File	Sealed → materialised	The record's own grade
1	NASA CLPS: Pricing the Launch Day	LA-003	3 Jan 2024 → 8 Jan 2024 (5 d)	HIT (sealed at 0.85; the same text's delay clause did not occur — preserved under the hash)
2	Lloyd's Market: Grading an Unconventional Forecast	the launch record entire	—	not one call but the ledger itself: 24 graded launch-day calls (22/1/1/0) as of 5 Aug 2026, self-scored, re-grading harness public
3	The Xi Weekend	IA-CN-2022	24 Sep 2022 → 27 Sep 2022 (3 d)	HIT at 0.95 — with the record's own 0.70 blind floor published (≈half a bit of yield)
4	Thirty-Six Hours in Moscow	IA-RU-005	6 Mar 2023 → 24 Jun 2023 (110 d)	HIT on class terms; breadth conceded in the file ("class only; unactionably broad as tasking")
5	The Tariff Shock	IA-MKT-002	2 Apr 2025 → 27 Jun 2025 (86 d)	HIT; the file concedes "the shock will pass" is the higher-base-rate side
6	The Interceptor Window	IA-RU-016	4 Sep 2023 → 11 Apr 2024 (220 d)	HIT on the dated window; the target inference belonged to the reader, never the seal
7	The Storm on the Frontage	IA-RU-006	4 Sep 2023 → 26 Nov 2023 (83 d)	HIT at 0.78, scenario-class only — mechanism mismatch and seasonal base rate conceded

8	The August Package	IA-US24-008	27 Aug 2024 → Sep–Nov 2024 (70 d)	PARTIAL (axis right, window wrong by three weeks) + a companion sub-claim MISS, both retained
9	The Swap-Equity Map	IA-RU-017	30 Apr 2023 → 1 Aug 2024 (460 d)	HIT on the named mechanism; the subject clause hedged and unbounded, conceded
10	The Trade-Profile Question	IA-RU-030	4 Sep 2023 → 22 Aug 2024 (353 d)	PARTIAL — two of four named vectors fired, two (grain, hijack) died and are marked as failures
11	Two Readers, One Input	LA-021 · LA-022	15 · 18 Apr 2026 → 18 · 19 Apr 2026 (3 d · 1 d)	LA-021's window ranking NEAR (never tested); LA-022's day call HIT at a 0.12 prior — with the Window-2 confirmation conceded as scheduling coincidence
12	The Founder Window	IA-SX-001	19 Mar 2025 → 5 Jun 2025 (78 d)	UNGRADED, by the record's own rule — no probability staked at seal, none backfilled, though the outcome would have flattered the aggregate
13	The Exit Polls	IA-IN24-001	10 Aug 2023 → 4 Jun 2024 (299 d)	HIT on the seat-ceiling call — inside a reading with 3 MISS (West Bengal cleanly) and 5 PARTIAL among ~32 sub-forecasts, all retained
14	The Substitute Port	IA-RU-019	4 Sep 2023 → 29 May 2024 (268 d)	HIT on campaign character and intent; the maximal "beyond repair" end-state NEVER occurred, carried as an explicit caveat
15	The Grind, Not the Gamble	IA-IC-2023	2 Nov 2023 → window closed empty	the dated attack DID NOT OCCUR; ungraded as prevention-class (scores in neither direction); the record refuses credit for the consensus-aligned posture read
16	The Structure That Survives Either Winner	IA-TR-2023	8 May 2023 → 28 May 2023+ (20 d)	headline winner MISSED (with the whole polling field); six structural elements resolved; ungraded on the Brier by rule — a hypothetical grading would COST the ledger, computed in the file
17	Sixty-One Days to the Fifteenth Ballot	IA-US22-001	7 Nov 2022 → 7 Jan 2023 (61 d)	HIT on all three clauses — with the record's own two-lead-times qualification (the no-wave half resolved in ~9 days)
18	The City Hall Window	IA-NYC-001	28 Oct 2025 → 4 Nov 2025 (7 d)	correct and DECLINED — the record refuses to grade its own consensus-confirming hit (~1.3 bits); the case exists because being right added nothing
19	The Day, Not the Dose	IA-USCV-001	14 Dec 2020 → ongoing (0 d lead)	UNGRADED by rule — zero-lead, non-falsifiable day-reading; floor honestly rated ~1-in-5; confound conceded as at least as plausible

20	The Two Days	LA-017	19 Feb 2026 → 21 Feb 2026 (2 d public lead)	tank-pressurization mechanism HIT (priors 0.10/0.15 published); the same map's telemetry mechanism NEAR — reality removed the test — and carried forward
21	The Third Shape	IA-RU-044	28 Feb 2022 → 30 Sep 2022 (214 d)	HIT at self-banded 0.90 — with the map-momentum objection, the unquantified "significant," and the untestable "without war" branch all conceded
22	The Festival Window	IA-IN24-NS03	16 Aug 2024 → 11 Sep 2024 (26 d)	HIT on the incident-spike class only; NO improbability figure computed ("festival-window frictions are not rare"); the same sitting carries 8 HIT · 2 NEAR · 1 MISS
23	The Do-Not-Spend Decision	IA-US24-010	29 Aug 2024 → 5 Nov 2024 (68 d)	HIT priced at almost nothing (1-in-4 floor, contrarian 1-of-5, no premium claimed) — shipped on the same clock as the run's GA/MI/NV MISSES
24	The Long File	IA-RU-001	28 Feb 2022 → 30 Apr 2024 milestone (792 d)	OPEN through 2035 — a tracking judgment; the base-rate objection (sanctions rarely achieve maximal objectives) conceded in the file
25	The Warning Delivered Twice	IA-RU-008	4 Sep 2023 → 22 Mar 2024 (200 d)	HIT at 8.71 bits (~1-in-420) — actionability scored honestly low (54); the counterfactual explicitly unclaimed; the political-dismissal confound conceded
26	Two Years Before the Freeze	IA-RU-011	26 Feb 2023 → 18 Mar 2025 (751 d)	HIT at 0.9 (~9.4 bits) — with the July-2023 open-statement deflation and the partial/sequential character of the freeze conceded
27	Scheduling Money Against a Window	IA-RU-042	26 May 2025 → 23 Nov 2025 (181 d)	the negotiation-window branch materialised on its own terms — an *opportunity*, no settlement; the same paragraph's May-9 branch is a retained MISS (IA-RU-022)
28	Litigation and Time	IA-INGEO-001	26 Oct 2023 → 12 Feb 2024 (109 d)	HIT at 0.78, priced modestly (~2.6 bits) — an established appellate off-ramp; no claim any government read or acted on it

29	The Central-Count Night	IA-US24-006	23 Aug 2024 → 6 Nov 2024 (74 d)	HIT (~1-in-512) on the count-trajectory + integrity-class clause; the overnight window conceded as statute, the tabulator specifics as outcome-only; GA/MI companion MISSES retained
30	The Membership Was Always the Electorate	IA-UK22-001	8 Aug 2022 → 5 Sep 2022 (28 d)	HIT at ~1-in-30 (~4.9 bits) — flagged **not** against-consensus in its own metadata; actionability 44 (no trading call in the text); the gilt crisis is aftermath, never foreseen
31	The Paralysis Assumption	IA-PS-2023	11 Nov 2023 → 25 Mar 2024 (135 d)	institutional turn in-window — **UNGRADED by the record's ethics policy** (a scoreable hit declined); no operational effect claimed for the resolution
32	Finishing the File Before the Seam	IA-DL-2025	18 Sep 2025 → open (windows to 5 Feb 2027)	**excluded from grading on ethics**; claims nothing about the subject's health or longevity; two windows remain ahead; treated as a planning deadline only
33	Fifty-Three Days Before Norris	IA-YS-2024	26 Mar 2024 → headline MISS	the dated eruption **missed outright** (no scientific basis for exact-time forecasting); the floor clause priced at ~1-in-150 with no credit for the failed date; excluded from the ledger as prevention-class
34	Planning on Continuity	IA-RU-045	19 Nov 2022 → 17 Feb 2024 (455 d)	HIT — with the momentum objection (record harvest in the bins) and the sanctions-carve-out objection conceded, and the same sitting's mixed resources clause printed
35	Holding the Line	IA-MKT-001	8 Feb 2023 → 24 May 2024 (471 d)	**NEAR** — three of four legs landed; the fourth cut by an exogenous indictment; live upgrade trigger unfired; the smart-money base-case objection conceded

36	The Four-Day Fuse	IA-RU-034	19 Nov 2022 → 23 Nov 2022 (4 d)	**PARTIAL** — fuse and friction landed; "successfully obstruct" failed, refuted by the target's own bulletins; two pre-seal discounts (elevated base rate; public fallback doctrine) conceded
37	A Named Actor, Sixteen Days Early	IA-RU-002	26 Feb 2023 → 14 Mar 2023 (16 d)	HIT on all rubric axes — with the coverage objection (one line in a 31-item thread) conceded; the EnvyScout mechanism outcome-only; ceiling is raised cost, never prevention
38	Planning Paris on the Neutral Pathway	IA-RU-039	4 Sep 2023 → 29 Sep 2023 (25 d)	HIT on the operative direction clause; two of three clauses priced at ~zero and graded neither; framing kept exact (partial suspension, never reinstatement)
39	Retasking Before the Anniversary	IA-RU-020	1 Jan 2026 → open	kinetic window **PARTIAL / UNRESOLVED** on public evidence; no test asserted; zero corroborating chatter and a prior-year antecedent both conceded; the attribution narrative never endorsed
40	The Black Sea Test	five-call cluster (IA-RU-004/006/019/023/030)	2023–24	4 HIT · 1 PARTIAL, **consistent with** the theatre-wide base rate, not better; two hits carry conceded gaps; the small-sample objection is the case's own strongest voice

Read the last column as a set and the volume's quiet argument appears: hits of every strength from 8.71 bits down to almost nothing; partials with named dead vectors; NEARs that exist because reality removed the test; an outright failed dated call taught as exactly that; a headline miss shared with an entire polling field; entries the record refused to score — some on their own interest, some on ethics; a file still open a decade out; a window that never resolved; and a five-call cluster deliberately shown to be *consistent with* its own base rate rather than better than it. A casebook drawn from a record that only ever graded itself HIT would deserve the classroom's suspicion. This one hands the classroom its concessions pre-written — which is precisely what makes the cases teachable.

THE RECORD, AS OF THIS BOOK

All cases cite the grading ledger at its snapshot of **2 August 2026** (integrity hash 542e298e…): **106 graded calls — 82 HIT, 10 NEAR, 5 PARTIAL, 9 MISS — mean Brier 0.0903**, every miss retained at full weight, permanently. Cases 1

and 2 additionally cite the launch-day subset (24 calls, Brier 0.0349) as of 5 August 2026, per their own exhibit dates. The record is a living dataset: by the time you teach these cases the counts will have moved, and the correct citation practice — modeled in every exhibit — is *as-of date plus ledger hash*, never a bare number. The current state is always at jyotishintelligence.com/grading-ledger.json.

The record's honest edge, printed here as it is printed there: on a small, high-confidence corpus, a naive base-rate forecaster roughly ties the aggregate Brier. The headline average is not the achievement; the specificity, the lead times, and the retained misses are. Grading is self-scored against pre-committed criteria; no independent adjudication exists yet; the published re-grading harness is the standing invitation to supply one.

HOW A STUDENT VERIFIES A CASE TONIGHT

1. **Anteriority.** Open the sealed artifact link in the case's Exhibit 2. The platform timestamp — infrastructure the operator does not control — establishes the text existed before the event.
2. **Integrity.** Each sealed claim is SHA-256-hashed into a manifest anchored to Bitcoin via OpenTimestamps. The verifier at jyotishintelligence.com/verify recomputes the digests; no trust in the author is required, or requested.
3. **Grading.** The advisory page publishes the grade, the rationale, the confidence at seal, and — where they exist — the concessions. The ledger JSON carries every entry; the regrade kit invites you to overturn any of them.
4. **Outcomes.** Every load-bearing outside-world fact in every case is cited to public primary or trade sources independent of the record's operator, because the author's interest is disclosed on page one of every case.

USING THE CASES

Distribute **case bodies only**; hold the teaching notes — the no-verdict seam is the pedagogy, and the Contents lists every case and every teaching note with its own folio so the student page range is a clean clip (see *For the Instructor*, which also carries the recommended pairings). Printable single-case handouts and the templates pack are hosted at jyotishintelligence.com/cases (QR + URL at the end of each case's exhibits) and deposited at Zenodo. The six postures and the Decision Package form these notes reference are developed in full in Volume IV; no case requires it.

Everything here is given away on the terms stated in the front matter: teach it, copy it, adapt it, ingest it — with the name attached, so the credit lands where it belongs.

“...but what would you have done, that day, with what was actually on the desk?”

THE FORECASTING PROTOCOL · VOLUME V · THE BUSINESS SCHOOL CASE STUDIES
SEALED BEFORE · GRADED AFTER · GIVEN AWAY ON PURPOSE
JYOTISHINTELLIGENCE.COM

The case ends at the decision point — the outcome, the grade, and the record's own concessions wait in the teaching note

A rumor no desk can verify. A window no forecast can see. A supplier whose founder is the event. Forty decision cases for the classroom, built on forecasts sealed before their events and graded after — hits, partials, nears, failures, and the calls the record refused to score. Every case ends where real decisions end: before the answer. Students argue the posture; the teaching note holds the grade; and every sealed artifact can be checked from a phone, tonight.

Don't trust the record — recompute it.

The verifier, dataset & full method are public —
jyotishintelligence.com · Zenodo DOI 10.5281/zenodo.20630257

Vijay Jyotish is a forecaster and researcher and the founder of JYOTINT —
the open, sealed-and-scored forecasting standard this book gives you.

ISBN 979-8-9969760-7-2